arXiv:1702.02686v1 [stat.ML] 9 Feb 2017

Rate Optimal Estimation and Confidence Intervals for High-dimensional Regression with Missing Covariates

Yining Wang[1], Jialei Wang[3], Sivaraman Balakrishnan[1,2], and Aarti Singh[1]

[1]Machine Learning Department, Carnegie Mellon University
[2]Department of Statistics, Carnegie Mellon University
[3]Department of Computer Science, University of Chicago

February 10, 2017

Abstract

We consider the problem of estimating and constructing component-wise confidence intervals of a sparse high-dimensional linear regression model when some covariates of the design matrix are missing completely at random. A variant of the Dantzig selector (Candes & Tao, 2007) is analyzed for estimating the regression model and a de-biasing argument is employed to construct component-wise confidence intervals under additional assumptions on the covariance of the design matrix. We also derive rates of convergence of the mean-square estimation error and the average confidence interval length, and show that the dependency over several model parameters (e.g., sparsity s, portion of observed covariates ρ_*, signal level $\|\beta_0\|_2$) are optimal in a minimax sense.

1 Introduction

High-dimensional regression has been an active topic of research in statistics and machine learning over the past 20 years (Tibshirani, 1996; Efron et al., 2004; Donoho, 2006; Candès et al., 2006; Fan & Li, 2001). Generally speaking, the high-dimensional estimation problem concerns the setting where the number of variables (or features) is on par with, or even far exceeds the number of observations (data points) available. To make the estimation problem well-defined, it is usually assumed that only a small portion of the variables are related to the response variable. A typical high-dimensional regression model is defined as

$$y = X\beta_0 + \varepsilon; \qquad X \in \mathbb{R}^{n\times p}, \;\; \varepsilon|X \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 I) \tag{1}$$

where β_0 is a p-dimensional sparse linear model to be estimated and ε is i.i.d. Gaussian noise with zero mean and variance σ_ε^2. The number of observations n is assumed to be smaller than the dimension p of each data point, while it is assumed that $s < n$ components of β_0 is non-zero. Popular estimators for Eq. (1) include the Lasso (Tibshirani, 1996), the SCAD (Fan & Li, 2001) and the Dantzig selector (Candes & Tao, 2007), whose asymptotic rates of convergence and

model selection properties are well understood (Zhao & Yu, 2006; Bach, 2008; Bickel et al., 2009; Wainwright, 2009)

In many statistical applications, however, the full design (data) matrix X is not fully observed and missing/corrupted entries are common. For example, in a data set that records characterization of $p = 5520$ genes for $n = 46$ patients with soft tissue tumors (Nielsen et al., 2002), a total of 6.7% entries are missing; in addition, 78.6% of the 5520 genes and *all* of the 46 patients have at least one missing covariate. Under such scenario, classical methods like list-deletion is no longer applicable; imputation based methods require additional assumptions on the data generative model and might lead to invalid confidence intervals because the noise of the imputed values is not taken into consideration.

In this paper, we consider the problem of estimating and building component-wise confidence intervals of β_0 without imputing an incomplete design X. Let $R_{ij} \in \{0, 1\}$ be indicator variables of whether X_{ij} is missing and define the observable zero-filled $n \times p$ matrix $\bar{X}$ as

$$\bar{X}_{ij} = R_{ij} X_{ij}, \qquad 1 \le i \le n, 1 \le j \le p.$$

Following Loh & Wainwright (2012a), we assume that each element X_{ij} is missing independently and completely at random, and impose a random design model over X where each row of X is sampled i.i.d. from a sub-Gaussian distribution with zero mean and covariance Σ_0. The main contributions of this paper are as follows:

Estimation with unknown Σ_0 We analyze the noisy Dantzig selector estimator and show that its average squared error of estimating β_0 (i.e., $\mathbb{E}\|\hat{\beta} - \beta_0\|_2^2$) depends quadratically on ρ, the probability of observing each X_{ij}, and linearly on $\|\beta_0\|_2^2$. The dependency over ρ is better than existing estimators (Loh & Wainwright, 2012a; Chen & Caramanis, 2013) under similar settings, which depend on the 4th power of ρ. It is also proved that the dependency over ρ^2 cannot be removed in the minimax sense. Our bounds are not directly comparable to (Rosenbaum & Tsybakov, 2010, 2013) that do not make random design assumptions, whose statistical rates depend on $\|\beta_0\|_1$ instead of $\|\beta_0\|_2$.

Estimation with known Σ_0 We analyze a variant of the noisy Dantzig selector and show that its averaged square error dependes *linearly* on ρ and $\|\beta_0\|_2^2$. This improves over existing estimators (Loh & Wainwright, 2012b) for the known Σ_0 setting, which depend quadratically on ρ. In addition, it is shown that under the identity covariance case $\Sigma_0 = I$ the dependency over both ρ and $\|\beta_0\|_2$ is minimax optimal, as well as the dependency over conventional quantities s, p and n.

Component-wise Confidence Intervals with unknown Σ_0 Under the additional assumption that Σ_0^{-1} is sparse, coordinate-wise confidence intervals of β_0 are constructed by *de-biasing* the noisy Dantzig selector. The constructed confidence intervals are conditioned on X, with randomness over both the missing pattern $\{R_{ij}\}$ and noise ε. Furthermore, under the identity covariance case $\Sigma_0 = I$ it is shown that the length of the confidence interval matches the minimax rate up to universal constants.

One important difference from existing de-biased sparse estimators (Javanmard & Montanari, 2014; Cai et al., 2014; Zhang & Zhang, 2014; van de Geer et al., 2014) is that when $\bar{X}$ contains missing values, the de-biasing matrix $\hat{\Theta}$ (defined in Eq. (7)) correlates with the sparse estimator $\hat{\beta}$ (cf. Eqs. (2,3)), and the limiting distribution of the de-biased estimator depends on unseen covariates

in X. We use a variant of the CLIME estimator (Cai et al., 2011) to resolve the correlation issue and propose a data-driven estimator for the limiting variance of the de-biased estimator.

Notations For a vector x, we use $\|x\|_p = \left(\sum_j |x_j|^p\right)^{1/p}$ to denote the p-norm of x. For a matrix A, we use $\|A\|_{L_p}$ to denote the operator p-norm of A; that is, $\|A\|_{L_p} = \sup_{x\neq 0} \|Ax\|_p/\|x\|_p$. We also write $\|A\|_\infty$ for the max norm of a matrix: $\|A\|_\infty = \max_{j,k} |A_{jk}|$. For a positive semi-definite matrix A, let $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ be the largest and smallest eigenvalues of A. We use $\mathbb{B}_p(M) = \{x : \|x\|_p \leq M\}$ to denote the centered ℓ_p ball of radius M.

1.1 Related work

Rosenbaum & Tsybakov (2010) proposed the MU-selector for high-dimensional regression under an error-in-variable model, where the design matrix X is observed with deterministic (adversarial) measurement error W that is bounded in matrix max norm. The estimator was further generalized to handle missing data in (Rosenbaum & Tsybakov, 2013) for which de-biasing the covariance estimator leads to improved error bounds. Optimization algorithms and minimax rates when W is Gaussian white noise are derived in (Belloni et al., 2016).

Loh & Wainwright (2012a) analyzed a gradient descent algorithm for optimizing a non-convex Lasso-type loss function and derived rates of convergence from both statistical and optimization perspectives. Their analysis shows a quadruple dependency over the observation/missing rate for the mean-square estimation error and requires an upper RE condition. Loh & Wainwright (2012b) derived lower bounds on the minimax rate, assuming identity covariance for the design points and bounded signal level $\|\beta_0\|_2$. The error lower bound depends linear in observation/missing rate (Loh & Wainwright, 2012b). A general analytical framework for non-convex optimization in high-dimensional regression problems is presented in (Loh & Wainwright, 2015). A similar rate of convergence was established in (Chen & Caramanis, 2013) for orthogonal matching pursuit (OMP) type estimators.

Datta & Zou (2015) proposed CocoLasso, a variant of Lasso for error-in-variable models where a covariance estimate $\hat{\Sigma}$ is projected onto a positive semi-definite cone so that the resulting optimization problem is convex. Both additive and multiplicative measurement error models were considered in (Datta & Zou, 2015) and corresponding rates of convergence were derived.

2 Rate-optimal estimation of β_0

2.1 Problem setup and assumptions

Throughout this paper we make the following assumptions:

(A1) *Homogenous Gaussian noise*: $\varepsilon \sim \mathcal{N}_n(0, \sigma_\varepsilon^2 I)$ for some $\sigma_\varepsilon < \infty$.

(A2) *Sub-Gaussian random design*: each row of X is sampled i.i.d. from some underlying sub-Gaussian distribution with covariance Σ_0 and sub-Gaussian parameter $\sigma_x < \infty$. Assume $0 < \lambda_{\min}(\Sigma_0) \leq \lambda_{\max}(\Sigma_0) < \infty$. For notational simplicity we drop Σ_0 and use $\lambda_{\min}, \lambda_{\max}$ instead in the rest of this paper.

(A3) *Missing completely at random*: R_{ij} are independent Rademacher variables with $\Pr[R_{ij} = 1] = \rho_j$ for some $\rho_1, \cdots, \rho_p \in (0, 1)$. Also assume $\{R_{ij}\}_{i,j} \perp\!\!\!\perp X, \varepsilon$ and $\rho_* = \min_{1 \le j \le p} \rho_j > 0$.

(A4) *Sparsity*: The support set $J_0 = \mathrm{supp}(\beta_0) = \{j : |\beta_{0j}| \neq 0\}$ satisfies $|J_0| \le s$ for some $s \ll n$.

(A1), (A3) and (A4) are standard assumptions for high-dimensional regression with missing data, and (A2) implies (with high probability) a deterministic Restricted Eigenvalue (RE) condition (Bickel et al., 2009) on the sample covariance of X. which leads to a $s \log p/n$ fast rate for estimating β_0.

2.2 The noisy Dantzig selector

Define $\tilde{X} \in \mathbb{R}^{n \times p}$ and $\tilde{\Sigma} \in \mathbb{R}^{p \times p}$ as

$$\tilde{X}_{ij} = \frac{R_{ij} X_{ij}}{\rho_j}, \quad \tilde{\Sigma} = \frac{1}{n} \tilde{X}^\top \tilde{X} - D\mathrm{diag}\left(\frac{1}{n} \tilde{X}^\top \tilde{X}\right),$$

where $D = \mathrm{diag}(1 - \rho_1, \cdots, 1 - \rho_p)$ is a known $p \times p$ diagonal matrix. It is a simple observation that, conditioned on X, $\mathbb{E}\tilde{X} = X$ and $\mathbb{E}\tilde{\Sigma} = \hat{\Sigma} = \frac{1}{n} X^\top X$. Define the noisy Dantzig selector as

$$\hat{\beta}_n \in \mathrm{argmin}_{\beta \in \mathbb{R}^p} \left\{ \|\beta\|_1 : \left\| \frac{1}{n} \tilde{X}^\top y - \tilde{\Sigma}\beta \right\|_\infty \le \tilde{\lambda}_n \right\}, \tag{2}$$

where $\tilde{\lambda}_n > 0$ are tuning parameters. Eq. (2) is a variant of the Dantzig selector Candes & Tao (2007) and is in principle similar to the MU-selector in Rosenbaum & Tsybakov (2010). Note that Eq. (2) is always a convex optimization problem (regardless of whether $\tilde{\Sigma}$ is positive semi-definite) and hence can be efficiently solved.

We also consider a variant of the noisy Dantzig selector under the idealized scenario where the population covariance Σ_0 for the design matrix is known. In particular, define $\check{\beta}_n$ as the solution of

$$\check{\beta}_n \in \mathrm{argmin}_{\beta \in \mathbb{R}^p} \left\{ \|\beta\|_1 : \left\| \frac{1}{n} \tilde{X}^\top y - \Sigma_0 \beta \right\|_\infty \le \check{\lambda}_n \right\}. \tag{3}$$

Note that the covariance estimate $\tilde{\Sigma}$ is replaced with the known population covariance Σ_0 in Eq. (3). The estimator $\check{\beta}_n$ is primarily for theoretical considerations, as Σ_0 is in general unknown in most applications. In following sections we show that $\check{\beta}_n$ achieves exact minimax rates for recovering β_0 with missing covariates.

2.3 Rates of convergence and minimax lower bounds

Theorem 2.1 establishes upper bounds on the mean square estimation error of β_0. Eq. (4) corresponds to the setting where the population covariance Σ_0 is known and Eq. (5) holds when Σ_0 is unknown.

Theorem 2.1. *Assume (A1) to (A4). If* $\frac{\log p}{\rho_*^2 n} \to 0$ *and* $\check{\lambda}_n \asymp (\sigma_x^2\|\beta_0\|_2 + \sigma_x\sigma_\varepsilon)\sqrt{\frac{\log p}{\rho_* n}}$*, then*

$$\|\check{\beta}_n - \beta_0\|_2 \le O_{\mathbb{P}}\left\{\frac{\sigma_x^2\sqrt{s}}{\lambda_{\min}}\left(\|\beta_0\|_2\sqrt{\frac{\log p}{\rho_* n}} + \frac{\sigma_\varepsilon}{\sigma_x}\sqrt{\frac{\log p}{\rho_* n}}\right)\right\}. \tag{4}$$

In addition, if $\max\left\{\frac{\sigma_x^4 s\log(\sigma_x p/\rho_*)}{\rho_*^3\lambda_{\min}^2 n}, \frac{\log p}{\rho_*^4 n}\right\} \to 0$ *and* $\tilde{\lambda}_n \asymp \sigma_x^2\|\beta_0\|_2\sqrt{\frac{\log p}{\rho_*^2 n}} + \sigma_x\sigma_\varepsilon\sqrt{\frac{\log p}{\rho_* n}}$*, then*

$$\|\hat{\beta}_n - \beta_0\|_2 \le O_{\mathbb{P}}\left\{\frac{\sigma_x^2\sqrt{s}}{\lambda_{\min}}\left(\|\beta_0\|_2\sqrt{\frac{\log p}{\rho_*^2 n}} + \frac{\sigma_\varepsilon}{\sigma_x}\sqrt{\frac{\log p}{\rho_* n}}\right)\right\}. \tag{5}$$

Finally, $\|\hat{\beta}_n - \beta_0\|_1 \le 2\sqrt{s}\|\hat{\beta}_n - \beta_0\|_2$ *and* $\|\check{\beta}_n - \beta_0\|_1 \le 2\sqrt{s}\|\check{\beta}_n - \beta_0\|_2$ *with probability* $1 - o(1)$.

Compared to Loh & Wainwright (2012a) our bounds are better by an $O(1/\rho^*)$ factor for $\hat{\beta}_n$ when Σ_0 is unknown and an $O(1/\rho_*^{3/2})$ factor better when Σ_0 is known. Our bounds are not directly comparable to Rosenbaum & Tsybakov (2010) that considers a fixed-design setting with no stochastic model assumed over X. We however remark that error bounds in Rosenbaum & Tsybakov (2010) depend on $\|\beta_0\|_1$, which could be a factor of $\sqrt{s}$ worse than $\|\beta_0\|_2$.

We next present minimax lower bounds for the L_2 estimation error $\|\hat{\beta}_n - \beta_0\|_2^2$. We first consider the simple case where the population covariance Σ_0 is the identity.

Theorem 2.2. *Suppose* $4 \le s < 4p/5$, $\frac{s\log(p/s)}{\rho_* n} \to 0$ *and* $\Sigma_0 = I$*. Then*

$$\begin{aligned}&\inf_{\hat{\beta}_n}\sup_{\beta_0 \in \mathbb{B}_2(M)\cap\mathbb{B}_0(s)} \mathbb{E}\|\hat{\beta}_n - \beta_0\|_2^2 \\ &\quad \ge C_0 \cdot \min\left\{\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2, e^{0.5c^2(1-\rho_*)s}\sigma_\varepsilon^2\right\} \cdot \min\left\{\sqrt{\frac{s\log(p/s)}{(1-\rho_*)^2 n}}, \frac{s\log(p/s)}{\rho_* n}\right\}. \end{aligned} \tag{6}$$

Here $C_0 > 0$ *is a universal constant and* $c > 0$ *is an arbitrary constant.*

Remark 2.1. Under the additional assumption that $\frac{(1-\rho_*)^2 s\log(p/s)}{\rho_*^2 n} \to 0$, the right-hand side of Eq. (6) can be simplified to

$$C_0 \cdot \min\left\{\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2, e^{0.5c^2(1-\rho_*)s}\sigma_\varepsilon^2\right\}\frac{s\log(p/s)}{\rho_* n}.$$

Suppose that $s/p \to 0$ and hence $\log p$ and $\log(p/s)$ are of the same order. If the missing rate $(1-\rho_*)$ is at least a constant and the sparsity level s or the noise level σ_ε is not too small, the term $e^{c^2(1-\rho_*)s}\sigma_\varepsilon^2$ is negligible because it increases exponentially with s. This term arises in the lower bound because on a subset of of the design points whose size decreases exponentially fast with s, the covariates corresponding to the support of β_0 are fully observed. Apart from this term, the lower bound matches the estimation error rate of $\check{\beta}_n$ when $\Sigma_0 = I$, corresponding to $\lambda_{\min} = \sigma_x = 1$.

The setting where Σ_0 is unknown is more complicated because of the $1/\rho_*^2$ dependency in the upper bound of $\hat{\beta}_n$ (of squared ℓ_2 estimation error). The following theorem shows that such dependency is unavoidable if the population covariance Σ_0 is unknown.

Theorem 2.3. *Let $\gamma_0 \in (0, 1/2)$ be an arbitrary small positive constant and suppose $s \geq 4$, $\max\{\frac{\sigma_\varepsilon^2}{M^2\rho_* n}, \frac{1}{\gamma_0 \rho_*^2 n}\} \to 0$. Define $\Lambda(\gamma_0) = \{\Sigma_0 \in \mathbb{S}_+^p : 1-\gamma_0 \leq \lambda_{\min}(\Sigma_0) \leq \lambda_{\max}(\Sigma_0) \leq 1+\gamma_0\}$, where $\mathbb{S}_+^p$ is the class of all positive definite $p \times p$ matrices. Then for any fixed $j \in \{1, \cdots, p\}$ it holds that*

$$\inf_{\hat{\beta}_n} \sup_{\substack{\beta_0 \in \mathbb{B}_2(M) \cap \mathbb{B}_0(s) \\ \Sigma_0 \in \Lambda(\gamma_0)}} \mathbb{E}|\hat{\beta}_{nj} - \beta_{0j}|^2 \geq C_1 \cdot \max\left\{\frac{\sigma_\varepsilon^2}{\rho_* n}, \min\left(\frac{1-\rho_*}{1+2c} M^2, e^{0.5c^2(1-\rho_*)s}\sigma_\varepsilon^2\right) \frac{1}{\rho_*^2 n}\right\},$$

where $C_1 > 0$ is a universal constant that only depends on γ_0 and $c > 0$ is an arbitrary constant.

3 Confidence intervals of regression coefficients

We describe a method that builds confidence intervals over the estimated $\hat{\beta}_n$ by de-biasing the noisy Dantzig selector. We need the following additional assumption to justify the proposed approach:

(A5) There exist $b_0, b_1 < \infty$ such that each row (and column) of Σ_0^{-1} belongs to $\mathbb{B}_0(b_0) \cap \mathbb{B}_1(b_1)$. That is, each row of Σ_0^{-1} is b_0-sparse and $\|\Sigma_0^{-1}\|_{L_1} \leq b_1$.

Condition (A5) allows the usage of CLIME or node-wise Lasso to estimate an approximate "inverse" of Σ_0 that asymptotically de-biases an estimate $\hat{\beta}_n$. Similar conditions for high-dimensional inference were studied in (van de Geer et al., 2014). We discuss potential settings where (A5) could be relaxed in Sec. 5.

3.1 The de-biased noisy Dantzig selector

Let $\hat{\Theta}$ be a $p \times p$ matrix obtained by solving the following optimization problem:

$$\hat{\Theta} \in \text{argmin}_{\Theta \in \mathbb{R}^{p\times p}} \left\{\|\Theta\|_1 : \|\tilde{\Sigma}\Theta - I_{p\times p}\|_\infty \leq \tilde{\nu}_n \text{ and } \|\Theta\tilde{\Sigma} - I_{p\times p}\|_\infty \leq \tilde{\nu}_n\right\}, \tag{7}$$

where $\tilde{\nu}_n > 0$ is some tuning parameter to be specified later. Eq. (7) is a missing data variant of the CLIME estimator proposed in Cai et al. (2011) for estimating precision matrices in high dimension. The following lemma formally establishes the performance of $\hat{\Theta}$ for estimating Σ_0^{-1}. Its proof is standard and for completeness we include it in the supplementary materials.

Lemma 3.1. *Under (A1), (A3) and (A5), suppose $\frac{\log p}{\rho_*^2 n} \to 0$ and $\tilde{\nu}_n \asymp \sigma_x^2 b_1 \sqrt{\frac{\log p}{\rho_*^2 n}}$. Then with probability $1 - o(1)$ it holds that $\max\{\|\hat{\Theta}\|_{L_1}, \|\hat{\Theta}\|_{L_\infty}\} \leq b_1$ and $\max\{\|\hat{\Theta} - \Sigma_0^{-1}\|_{L_1}, \|\hat{\Theta} - \Sigma_0^{-1}\|_{L_\infty}\} \leq 2\tilde{\nu}_n b_0 b_1$.*

With an estimator $\hat{\beta}_n$ of β_0 and $\hat{\Theta}$ obtained from solving Eq. (7), the de-biased estimator $\hat{\beta}_n^u$ is defined as

$$\hat{\beta}_n^u = \hat{\beta}_n + \hat{\Theta}\left(\frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\hat{\beta}_n\right). \tag{8}$$

We then have the following theorem that derives the limiting variance of $\hat{\beta}_n^u$.

Theorem 3.1. *Define* $\hat{\Gamma} \in \mathbb{R}^{p\times p}$ *as*

$$\hat{\Gamma} = \frac{\sigma_\varepsilon^2}{n} X^\top X + \frac{\sigma_\varepsilon^2}{n} \tilde{D}\mathrm{diag}(X^\top X) + \hat{\Upsilon},$$

where $\tilde{D} = \mathrm{diag}(\frac{1}{\rho_1} - 1, \cdots, \frac{1}{\rho_p} - 1)$ *and*

$$\hat{\Upsilon}_{jk} = \begin{cases} \frac{1}{n}\sum_{i=1}^n \sum_{t\neq j} \frac{1-\rho_t}{\rho_j \rho_t} X_{ij}^2 X_{it}^2 \beta_{0t}^2, & j = k; \\ \frac{1}{n}\sum_{i=1}^n \sum_{t\neq j,k} \frac{1-\rho_t}{\rho_t} X_{ij} X_{ik} X_{it}^2 \beta_{0t}^2, & j \neq k. \end{cases}$$

If the conclusion in Lemma 3.1 holds and

$$\sigma_x^2 b_0 b_1 \tilde{\nu}_n \left(\frac{\sigma_\varepsilon}{\sigma_x} \sqrt{\frac{\log p}{\rho_*}} + \|\beta_0\|_2 \sqrt{\frac{\log p}{\rho_*^2}} + \frac{\sqrt{n}\|\hat{\beta}_n - \beta_0\|_1}{\sigma_x^2 b_0 b_1} \right) \xrightarrow{p} 0, \tag{9}$$

then for any variable subset $S \subseteq [p]$ *with constant size it holds that with probability* $1 - o(1)$ *over the random design* X*,*

$$\sqrt{n}\left(\hat{\beta}_n^u - \beta_0\right)_{SS} \xrightarrow{d} \mathcal{N}_{|S|}\left(0, \left[\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1}\right]_{SS}\right) \quad \textit{conditioned on } X.$$

Remark 3.1. When the noisy Dantzig selector Eq. (2) is used for the initial estimation $\hat{\beta}_n$ and both $\tilde{\nu}_n$ and $\tilde{\lambda}_n$ are chosen at the rates specified in Theorem 2.1 and Lemma 3.1, then a sufficient condition for Eq. (9) to hold is that

$$\sigma_x^4 b_0 b_1^2 \sqrt{\frac{\log^2 p}{\rho_*^4 n}} \left(\frac{\sigma_\varepsilon \sqrt{\rho_*}}{\sigma_x} + \|\beta_0\|_2 \right) \left(1 + \frac{s}{\lambda_{\min} b_0 b_1} \right) \to 0. \tag{10}$$

Remark 3.2. We derive the asymptotic variance of $\sqrt{n}(\hat{\beta}_{nj}^u - \beta_{0j})$ for a specific coordinate j under the identity covariance setting $\Sigma_0 = I$ and demonstrate its rate optimality. For simplicity we also assume uniform observation rates $\rho_1 = \rho_2 = \cdots = \rho_p = \rho_*$. Fix j as a single coordinate and let $V_j = \mathbb{V}_{jj}\sqrt{n}(\hat{\beta}_n^u - \beta_0)$. By Theorem 3.1, when n is sufficiently large

$$V_j \xrightarrow{p} \hat{\Gamma}_{jj} \xrightarrow{p} \frac{\sigma_\varepsilon^2}{\rho_*} + \frac{1-\rho_*}{\rho_*^2} \sum_{t\neq j} \beta_{0t}^2 \leq \frac{\sigma_\varepsilon^2}{\rho_*} + \frac{1-\rho_*}{\rho_*^2} \|\beta_0\|_2^2. \tag{11}$$

Comparing Eq. (11) with Theorem 2.3, the variance V_j achieves the minimax rates of coordinate-wise estimation up to problem independent constants. Formally, under the additional assumption $\sigma_\varepsilon^2 \gg e^{-0.5c^2(1-\rho_*)s}\|\beta_0\|_2^2$ that σ_ε is not exponentially small, we have that

$$\limsup_{p,n\to\infty} \frac{V_j^2}{\inf_{\hat{\beta}'_n} \sup_{\beta'_0 \in \mathbb{B}_2(\|\beta_0\|_2)\cap\mathbb{B}_0(s), \Sigma'_0\in\Lambda(\gamma_0)} n\mathbb{E}_{0'}|\hat{\beta}'_{nj} - \beta'_{0j}|^2} \leq 2C_1^{-1}(1+2c),$$

where $C_1 > 0$ is a universal constant in Theorem 2.3.

3.2 Data-driven approximation of the limiting covariance

The limiting variance $\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1}$ derived in Theorem 3.1 depends on the full design matrix X and the population precision matrix Σ_0^{-1}, both of which are unavailable in practice when some covariates of X are missing. To overcome this difficulty, in this section a data-driven approximation of $\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1}$ is proposed, which only depends on the observed data $\tilde{X}$ and y.

Define $\tilde{\Gamma} = \frac{\sigma_\varepsilon^2}{n}\tilde{X}^\top\tilde{X} + \tilde{\Upsilon}$, where

$$\tilde{\Upsilon}_{jk} = \frac{1}{n}\sum_{i=1}^{n}\sum_{t\neq j,k}(1-\rho_t)\tilde{X}_{ij}\tilde{X}_{ik}\tilde{X}_{it}^2\hat{\beta}_{nt}^2,$$

for $j,k \in \{1,\cdots,p\}$. The following theorem shows that $\hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top$ is a good approximation of $\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1}$ when n is sufficiently large:

Theorem 3.2. *Suppose the conclusion in Lemma 3.1 holds,* $\frac{\log p}{\rho_*^4 n} \to 0$ *and* $\|\hat{\beta}_n - \beta_0\|_2 \xrightarrow{p} 0$. *Then*

$$\left\|\hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top - \Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1}\right\|_\infty \leq O_{\mathbb{P}}\left(\frac{\sigma_x^4 b_1^2\log^2 p}{\rho_*^2}\left\{\left(\|\beta_0\|_2^2 + \frac{\rho_*\sigma_\varepsilon^2}{\sigma_x^2}\right)\left(b_0\tilde{\nu}_n + \sqrt{\frac{\log p}{\rho_* n}}\right) + \|\beta_0\|_2\|\hat{\beta}_n - \beta_0\|_1\right\}\right).$$

Based on Theorems 3.1 and 3.2, an asymptotic $(1-\alpha)$ confidence interval of β_{0j} can be computed as

$$\mathrm{CI}_j(\alpha) = \left[\hat{\beta}_{nj}^u - \frac{\Phi^{-1}(1-\alpha/2)\sqrt{(\hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top)_{jj}}}{\sqrt{n}}, \hat{\beta}_{nj}^u + \frac{\Phi^{-1}(1-\alpha/2)\sqrt{(\hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top)_{jj}}}{\sqrt{n}}\right], \qquad (12)$$

where $\Phi^{-1}(\cdot)$ is the inverse function of the CDF of the standard Gaussian distribution.

4 Simulation results

4.1 Synthetic data

We fix $\sigma_\varepsilon = 0.1$ and set $\Sigma_0 = \Omega^{-1}$ where Ω is chosen to be the following banded matrix:

$$\Omega_{ij} = \begin{cases} 0.5^{|i-j|} & \text{if } |i-j| \leq 5 \\ 0 & \text{otherwise} \end{cases}.$$

Assume uniform observation rate $\rho_1 = \cdots = \rho_p = \rho$, which ranges from 0.5 to 0.9. The support set $J_0 \subset [p]$ of β_0 is selected uniformly at random, with $|J_0| = 10$. β_0 is then generated as $\beta_{0j} \overset{i.i.d.}{\sim} \text{Bernoulli}(+1,-1)$ for $j \in J_0$ and $\beta_{0j} = 0$ for $j \notin J_0$. Both the noisy Dantzig selector Eq. (2) and the noisy CLIME estimator Eq. (7) are solved using alternating direction methods of multipliers (ADMM).

4.1.1 Verification of asymptotic normality

We run 1000 independent realizations of $\{R, \bar{X}, y\}$ and study the distributions of $\sqrt{n}(\hat{\beta}_n^u - \beta_0)$. We plot the empirical distribution of

$$\hat{\delta}_j = \frac{\sqrt{n}(\hat{\beta}_{nj}^u - \beta_{0j})}{\sqrt{(\hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top)_{jj}}}$$

against the standard Normal distribution. Figure 1 shows that the empirical distribution of $\hat{\delta}_j$ agrees quite well with $\mathcal{N}(0,1)$. In addition, more observations (n) are required to deliver asymptotic normality when observation rates are low (e.g., $\rho = 0.5$).

4.1.2 Average CI coverage and length

We calculate the average coverage and length of the constructed confidence intervals from T independent realizations, defined as

$$\mathrm{Avgcov}(j) = \frac{1}{T}\sum_{i=1}^{T} \mathbb{I}(\beta_{0j} \in \mathrm{CI}_j^{(i)}(\alpha)), \quad \text{and} \quad \mathrm{Avglen}(j) = \frac{1}{T}\sum_{i=1}^{T} \mathrm{length}(\mathrm{CI}_j^{(i)}(\alpha)),$$

where $\mathrm{CI}_j(\alpha)$ is defined in Eq. (12). We also report the average coverage and length of coordinate-wise confidence intervals across a coordinate subset $J \subseteq [p]$, defined as

$$\mathrm{Avgcov}(J) = \frac{1}{|J|}\sum_{j\in J} \mathrm{Avgcov(j)} \qquad \text{and} \qquad \mathrm{Avglen}(J) = \frac{1}{|J|}\sum_{j\in J} \mathrm{Avglen(j)}.$$

Tables 1 summarize the results for various (n, p, ρ) settings.

4.2 Real data

In this section we conduct experiments on two datasets: DNA and Madelon[1], where the distribution of the design matrices are not necessarily sampled sub-Gaussian distributions. The DNA data contains 2000 instances and 180 covariates, while Madelon contains 2000 data points and 500 covariates. For these two datasets, we only uses their data matrix X and construct the response y according to the sparse linear regression model as specified in previous section. Following the simulation study, we randomly remove the observed covariates with probability $1 - \rho$, and then perform statistical inference based on the datasets with missing covariates. The performance of the constructed confidence intervals are reported in Table 3. We see that the proposed procedure can produced roughly normal estimates for the parameters of interest when ρ is not too small, this demonstrate that the method can be robust to violations of the statistical assumptions on the design matrix.

5 Discussion

Quadratic dependency over ρ_* Theorem 2.3 shows that if the population covariance Σ_0 of the design matrix X is unknown, then mean square estimation error of a particular component in β_0

[1] Available from https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/

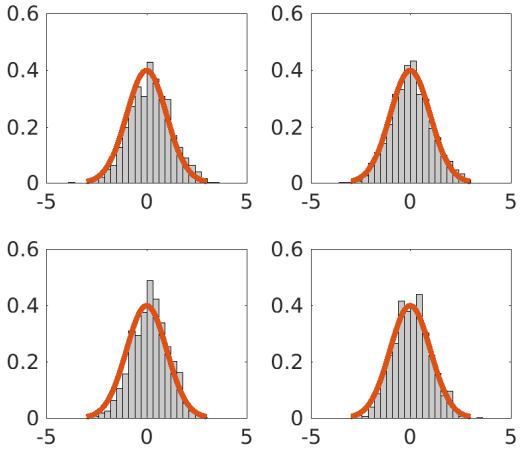

$n = 1500, p = 500, \rho = 0.9$

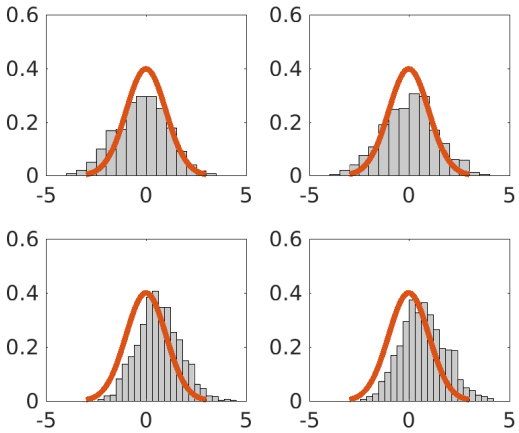

$n = 5000, p = 500, \rho = 0.7$

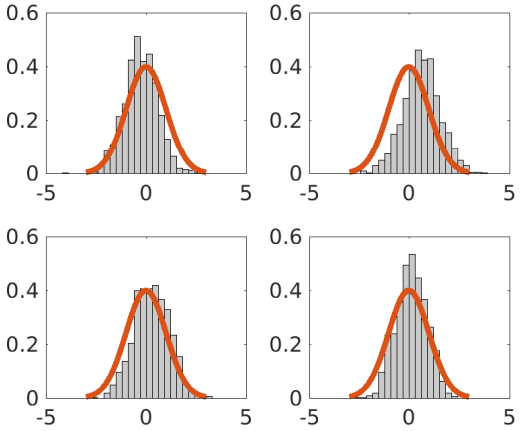

$n = 8000, p = 500, \rho = 0.5$

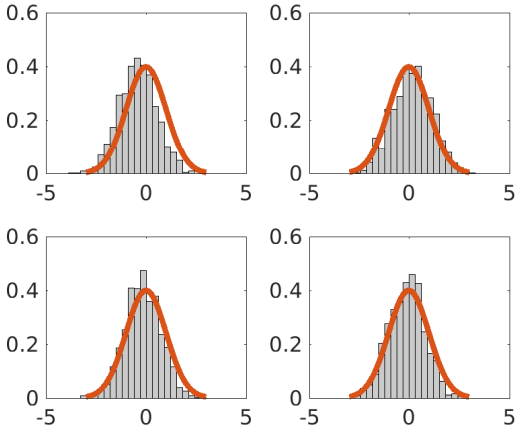

$n = 12000, p = 500, \rho = 0.5$

Figure 1: Empirical distribution and density of $\hat{\delta}_j = \frac{\sqrt{n}(\hat{\beta}^u_{nj} - \beta_{0j})}{\sqrt{(\hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top)_{jj}}}$ of 1000 independent realizations. Top row in each subfigure: two coordinates randomly chosen from J_0; bottom row in each subfigure: two coordinates randomly chosen from J_0^c. Red curve: density of the standard Normal distribution.

Table 1: 95% confidence intervals for high-dimensional regression with missing data when $\rho \in [0.7, 0.9]$.

(n, p, ρ)	Random $j \in J_0$		Random $j \notin J_0$		J_0		J_0^c	
	Avgcov	Avglen	Avgcov	Avglen	Avgcov	Avglen	Avgcov	Avglen
(1000,200,0.9)	0.941	0.182	0.951	0.192	0.938	0.208	0.966	0.187
(1000,200,0.8)	0.945	0.318	0.948	0.329	0.944	0.334	0.979	0.331
(1000,200,0.7)	0.952	0.494	0.983	0.540	0.949	0.547	0.989	0.529
(1500,500,0.9)	0.931	0.155	0.966	0.170	0.945	0.183	0.971	0.158
(1500,500,0.8)	0.927	0.278	0.982	0.294	0.937	0.308	0.985	0.284
(1500,500,0.7)	0.963	0.415	0.994	0.469	0.971	0.497	0.995	0.450
(2000,1000,0.9)	0.947	0.144	0.974	0.144	0.949	0.160	0.975	0.139
(2000,1000,0.8)	0.967	0.249	0.987	0.264	0.939	0.281	0.990	0.254
(2000,1000,0.7)	0.952	0.378	0.995	0.422	0.930	0.451	0.997	0.409
(3000,2000,0.9)	0.958	0.116	0.954	0.118	0.951	0.133	0.981	0.115
(3000,2000,0.8)	0.919	0.202	0.979	0.220	0.948	0.236	0.993	0.212
(3000,2000,0.7)	0.891	0.315	0.998	0.349	0.950	0.372	0.998	0.348

Table 2: 95% confidence intervals for regression with missing data when $\rho = 0.5$.

(n, p, ρ)	Random $j \in J_0$		Random $j \notin J_0$		J_0		J_0^c	
	Avgcov	Avglen	Avgcov	Avglen	Avgcov	Avglen	Avgcov	Avglen
(1000,200,0.5)	0.928	1.051	0.998	1.223	0.942	1.384	0.999	1.194
(2000,200,0.5)	0.971	0.715	0.997	0.849	0.971	0.799	0.995	0.813
(3000,200,0.5)	0.956	0.574	0.976	0.644	0.961	0.668	0.989	0.640
(4000,200,0.5)	0.936	0.468	0.984	0.541	0.943	0.527	0.986	0.534
(1500,500,0.5)	0.986	0.795	0.978	0.911	0.756	0.954	1.000	0.896
(3000,500,0.5)	0.849	0.510	0.899	0.575	0.479	0.634	0.998	0.572
(8000,500,0.5)	0.972	0.352	0.978	0.408	0.908	0.417	0.988	0.403
(12000,500,0.5)	0.941	0.272	0.965	0.315	0.936	0.328	0.976	0.309

Table 3: 95% confidence intervals for regression with missing data on real world datasets.

$(\text{dataset}, \rho)$	Random $j \in J_0$		Random $j \notin J_0$		J_0		J_0^c	
	Avgcov	Avglen	Avgcov	Avglen	Avgcov	Avglen	Avgcov	Avglen
(DNA,0.9)	0.924	0.120	0.956	0.128	0.937	0.128	0.957	0.129
(DNA,0.8)	0.908	0.195	0.959	0.216	0.926	0.212	0.965	0.218
(DNA,0.7)	0.888	0.286	0.967	0.318	0.925	0.314	0.973	0.317
(DNA,0.5)	0.713	0.464	0.964	0.516	0.745	0.512	0.976	0.519
(Madelon,0.9)	0.943	0.095	0.963	0.101	0.949	0.098	0.945	0.105
(Madelon,0.8)	0.966	0.167	0.976	0.174	0.961	0.181	0.971	0.223
(Madelon,0.7)	0.962	0.229	0.977	0.236	0.956	0.253	0.977	0.261
(Madelon,0.5)	0.663	0.334	0.977	0.357	0.682	0.377	0.965	0.356

must depend quadratically on the observation ratio ρ_*. We conjecture that such results also hold for the estimation error of the entire regression model β_0 as well. More specifically, we conjecture that under suitable finite-sample conditions,

$$\inf_{\hat{\beta}_n} \sup_{\substack{\beta_0 \in \mathbb{B}_2(M) \cap \mathbb{B}_0(s) \\ \Sigma_0 \in \Lambda(\gamma_0)}} \mathbb{E}\|\hat{\beta}_n - \beta_0\|_2^2 \geq C_1' \cdot \max\left\{ \frac{\sigma_\varepsilon^2 s \log p}{\rho_* n}, \min\left(\frac{1-\rho_*}{1+2c} M^2, e^{c^2(1-\rho_*)s} \sigma_\varepsilon^2 \right) \frac{s \log p}{\rho_*^2 n} \right\}.$$

We are, however, unable to generalize our construction of diffcult cases in the proof of Theorem 2.3 (cf. Sec. 6.4) to handle $\mathbb{E}\|\hat{\beta} - \beta_0\|_2^2$. The current construction relies on a carefully designed set of covariance matrices that leak no information unless both X_1 and X_j are observed. Extending such construction to multiple covariates requires new ideas.

Finite-sample conditions The finite-sample condition (i.e., relationship between n and the other model parameters for the asymptotic error bound to hold) required for Eq. (5) in Theorem 2.1 is slightly more restrictive than the actual error bound suggests. The condition arises from the use of Bernstein-type concentration inequalities, where the variance of a concentrated empirical sum is much smaller than its high-order moments. We are not sure whether such finite-sample conditions are results of a fundamental information-theoretical limitation, or can be avoided by a more refined analysis.

Confidence intervals constructed in Sec 3 requires more stringent conditions to be asymptotically level α. Eq. (10) suggests that at least $n \gg s^2 \rho_*^{-4} \log p$ needs to be satisfied. On the other hand, Cai & Guo (2015) shows that no adaptive confidence intervals exist under the regime of $n < s^2$.

On Condition (A5) (A5) requires the population precision matrix Σ_0^{-1} to be sparse, which could be restrictive as the preicision matrix is only a nuisance parameter and confidence intervals of β_0 do not necessarily require good estimation of Σ_0^{-1}. Javanmard & Montanari (2014) drops such sparse precision conditions at the cost of asymptotic efficiency of the average length of the resulting CI.

However, the techniques in (Javanmard & Montanari, 2014) cannot be easily adapted to the missing data case because both estimates $\hat{\Theta}$ and $\hat{\beta}_n$ depend on the randomness of the missing patterns R. In our proof we circumvent this issue by connecting to the deterministic population precision Σ_0^{-1}, which we do not know how to generalize to the case when Σ_0^{-1} is not sparse.

6 Proofs

6.1 Additional notations on concentration bounds

Definition 6.1. *Let A, B be random or deterministic square matrices of the same size and ε be a random vector of i.i.d. $\mathcal{N}(0, \sigma_\varepsilon^2)$ components. Let $\varphi_{u,v}(A, B; \log N)$, $\varphi_{u,\infty}(A, B; \log N)$, $\varphi_{\varepsilon,\infty}(A)$ be terms such that, with probability $1 - o(1)$, for all subset $\mathcal{S}$ of vectors with $|\mathcal{S}| \le N$, the following hold for all $u, v \in \mathcal{S}$:*

$$\begin{aligned} \left|u^\top (A - B) v\right| &\le \varphi_{u,v}(A, B; \log N) \cdot \|u\|_2 \|v\|_2; \\ \left\|A^\top \varepsilon\right\|_\infty &\le \varphi_{\varepsilon,\infty}(A) \cdot \sigma_\varepsilon. \end{aligned}$$

Note that $\varphi_{u,v}(\cdot, \cdot)$ is symmetric and satisfies the triangle inequality. Also, infinity norms like $\|A - B\|_\infty$ or $\|(A - B)u\|_\infty$ for a fixed u can be upper bounded by $\varphi_{u,v}(A, B; O(\log \dim(A)))$, by considering the set of unit vectors $\{e_1, \cdots, e_{\dim(A)}\}$.

6.2 Proof of Theorem 2.1

We need the following two concentration lemmas, which are proved in the supplementary material.

Lemma 6.1. *Denote random matrices $A^{(\ell)}$, $\ell \in \{0, 1, 2\}$ as $A^{(0)} = \hat{\Sigma}$, $A^{(1)} = \frac{1}{n}\tilde{X}^\top X$ and $A^{(2)} = \tilde{\Sigma}$, respectively. Then for $\ell \in \{0, 1, 2\}$:*

$$\varphi_{u,v}\left(A^{(\ell)}, \Sigma_0; \log N\right) \le O\left(\sigma_x^2 \max\left\{\frac{\log N}{\rho_*^{1.5\ell} n}, \sqrt{\frac{\log N}{\rho_*^{\ell} n}}\right\}\right).$$

Lemma 6.2. *If $\frac{\log p}{\rho_* n} \to 0$ then $\varphi_{\varepsilon,\infty}(\frac{1}{n}\tilde{X}) \le O(\sigma_\varepsilon \sigma_x \sqrt{\frac{\log p}{\rho_* n}})$.*

We present the following lemma commonly known as the *basic inequality* in high-dimensional inference literature. Its proof is given in the supplementary material.

Lemma 6.3 (Basic inequality)**.** *Suppose $\frac{\log p}{\rho_*^4 n} \to 0$ for $\hat{\beta}_n$ or $\frac{\log p}{\rho_*^2 n} \to 0$ for $\check{\beta}_n$, and let $J_0 = \mathrm{supp}(\beta_0)$ be the support of β_0. If $\tilde{\lambda}_n \ge \Omega\{\sigma_x \sqrt{\frac{\log p}{n}}(\frac{\sigma_x \|\beta_0\|_2}{\rho_*} + \frac{\sigma_\varepsilon}{\sqrt{\rho_*}})\}$ and $\check{\lambda}_n \ge \Omega\{\sigma_x \sqrt{\frac{\log p}{\rho_* n}}(\sigma_x \|\beta_0\|_2 + \sigma_\varepsilon)\}$, then with probability $1 - o(1)$ we have that $\|(\hat{\beta}_n - \beta_0)_{J_0^c}\|_1 \le \|(\hat{\beta}_n - \beta_0)_{J_0}\|_1$ and $\|(\check{\beta}_n - \beta_0)_{J_0^c}\|_1 \le \|(\check{\beta}_n - \beta_0)_{J_0}\|_1$.*

Definition 6.2 (Restricted eigenvalue condition)**.** *A $p \times p$ matrix A is said to satisfy* $\mathrm{RE}(s, \phi_{\min})$ *if for all $J \subseteq [p]$, $|J| \le s$ the following holds:*

$$\inf_{h \neq 0, \|h_{J^c}\|_1 \le \|h_J\|_1} \frac{h^\top A h}{h^\top h} \;\ge\; \phi_{\min}.$$

The following lemma is proved in the supplementary material.

Lemma 6.4. *Suppose* $\frac{\sigma_x^4 s \log(\sigma_x \log p/\rho_*)}{\rho_*^3 \lambda_{\min}^2 n} \to 0$. *Then with probability* $1 - o(1)$, $\tilde{\Sigma}$ *satisfies* $\mathrm{RE}(s, (1 - o(1))\lambda_{\min}(\Sigma_0))$.

We are now ready to prove Theorem 2.1 that establishes the rate of convergence of the noisy Dantzig selector estimators. We consider $\hat{\beta}_n$ first. Define $\tilde{\lambda}_n \mu = \frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\hat{\beta}_n$. By $y = X\beta_0 + \varepsilon$, we have that

$$\tilde{\Sigma}(\hat{\beta}_n - \beta_0) = \left(\frac{1}{n}\tilde{X}^\top X - \Sigma_0\right)\beta_0 + \left(\Sigma_0 - \tilde{\Sigma}\right)\beta_0 - \tilde{\lambda}_n \mu + \frac{1}{n}\tilde{X}^\top \varepsilon.$$

Multiply both sides by $(\hat{\beta}_n - \beta_0)$ and apply Hölder's inequality:

$$\begin{aligned}
&(\hat{\beta}_n - \beta_0)^\top \tilde{\Sigma}(\hat{\beta}_n - \beta_0) \\
&\le \|\hat{\beta}_n - \beta_0\|_1 \left\{ \left\| \left(\frac{1}{n}\tilde{X}^\top X - \Sigma_0\right)\beta_0 \right\|_\infty + \left\| \left(\Sigma_0 - \tilde{\Sigma}\right)\beta_0 \right\|_\infty + \tilde{\lambda}_n \|\mu\|_\infty + \left\| \frac{1}{n}\tilde{X}^\top \varepsilon \right\|_\infty \right\} \\
&\le \|\hat{\beta}_n - \beta_0\|_1 \cdot O_{\mathbb{P}} \left\{ \varphi_{u,v}\left(\frac{1}{n}\tilde{X}^\top X, \Sigma_0; \log p\right) \|\beta_0\|_2 + \varphi_{u,v}\left(\tilde{\Sigma}, \Sigma_0; \log p\right)\|\beta_0\|_2 + \tilde{\lambda}_n + \varphi_{\varepsilon,\infty}\left(\frac{1}{n}\tilde{X}\right)\sigma_\varepsilon \right\} \\
&\le \|\hat{\beta}_n - \beta_0\|_1 \cdot O_{\mathbb{P}} \left\{ \sigma_x^2 \|\beta_0\|_2 \sqrt{\frac{\log p}{\rho_*^2 n}} + \sigma_x \sigma_\varepsilon \sqrt{\frac{\log p}{\rho_* n}} + \tilde{\lambda}_n \right\}.
\end{aligned}$$

Here the last inequality is due to Lemmas 6.1 and 6.2. Suppose $\frac{\sigma_x^4 s \log(\sigma_x p/\rho_*)}{\rho_*^4 \lambda_{\min}^2 n} \to 0$ and $\tilde{\lambda}_n$ is appropriately set as in Lemma 6.3. We then have

$$\|\hat{\beta}_n - \beta_0\|_1 \le 2\|(\hat{\beta}_n - \beta_0)_{J_0}\|_1 \le 2\sqrt{s}\|\hat{\beta}_n - \beta_0\|_2$$

by Lemma 6.3 and

$$(\hat{\beta}_n - \beta_0)^\top \tilde{\Sigma}(\hat{\beta}_n - \beta_0) \ge (1 - o(1))\lambda_{\min}\|\hat{\beta}_n - \beta_0\|_2^2$$

by Lemma 6.4. Chaining all inequalities we get

$$\begin{aligned}
\|\hat{\beta}_n - \beta_0\|_2 &\le O_{\mathbb{P}}\left(\frac{\sqrt{s}}{\lambda_{\min}} \left\{ \sigma_x^2 \|\beta_0\|_2 \sqrt{\frac{\log p}{\rho_*^2 n}} + \sigma_x \sigma_\varepsilon \sqrt{\frac{\log p}{\rho_* n}} + \tilde{\lambda}_n \right\} \right) \\
&\le O_{\mathbb{P}}\left(\frac{\sqrt{s}}{\lambda_{\min}} \left\{ \sigma_x^2 \|\beta_0\|_2 \sqrt{\frac{\log p}{\rho_*^2 n}} + \sigma_x \sigma_\varepsilon \sqrt{\frac{\log p}{\rho_* n}} \right\} \right).
\end{aligned}$$

The ℓ_1 norm error bound $\|\hat{\beta}_n - \beta_0\|_1$ can be easily obtained by the fact that $\|\hat{\beta}_n - \beta_0\|_1 \le 2\sqrt{s}\|\hat{\beta}_n - \beta_0\|_2$.

Finally, consider $\check{\mu}_n$ and define $\check{\lambda}_n\check{\mu} = \frac{1}{n}\tilde{X}^\top y - \Sigma_0\check{\beta}_n$. Note that $\|\check{\delta}\|_\infty \le 1$ and

$$\Sigma_0(\check{\beta}_n - \beta_0) = \left(\frac{1}{n}\tilde{X}^\top X - \Sigma_0\right)\beta_0 - \check{\lambda}_n\check{\mu} + \frac{1}{n}\tilde{X}^\top\varepsilon.$$

Note in addition that $(\check{\beta}_n-\beta_0)^\top\Sigma_0(\check{\beta}_n-\beta_0) \ge \lambda_{\min}\|\check{\beta}_n-\beta_0\|_2^2$ by Assumption (A3). Subsequently, the same line of argument as $\hat{\beta}_n$ yields

$$\begin{aligned}\|\check{\beta}_n - \beta_0\|_2 &\le \frac{2\sqrt{s}}{\lambda_{\min}} \cdot O_{\mathbb{P}}\left\{\varphi_{u,v}\left(\frac{1}{n}\tilde{X}^\top X, \Sigma_0; \log p\right)\|\beta_0\|_2 + \tilde{\lambda}_n + \varphi_{\varepsilon,\infty}\left(\frac{1}{n}\tilde{X}\right)\sigma_\varepsilon\right\} \\ &\le O_{\mathbb{P}}\left\{\left(\sigma_x^2\|\beta_0\|_2 + \sigma_x\sigma_\varepsilon\right)\sqrt{\frac{s\log p}{\lambda_{\min}^2\rho_* n}}\right\}.\end{aligned}$$

6.3 Proof of Theorem 2.2

We consider the worst case with equal observation rates across covariates: $\rho_1 = \cdots = \rho_p = \rho_*$ and use Fano's inequality (Lemma A.1) to establish the minimax lower bound in Theorem 2.2. Construct hypothesis β as

$$\beta = (\underbrace{a, \cdots, a}_{\text{repeat } s/2 \text{ times}}, \underbrace{0, \pm\delta, 0, \cdots, \pm\delta, 0}_{\text{exactly } s/2 \text{ copies of } \delta}), \tag{13}$$

where $\delta \to 0$ is some parameter to be chosen later and $a = \sqrt{\frac{2M^2}{s} - \delta^2}$ is carefully chosen so that $\|\beta\|_2 = M$. Clearly $\beta \in \mathbb{B}_2(M) \cap \mathbb{B}_0(s)$. Let $d_H(\beta, \beta') = \sum_{j=1}^p I[\beta_j \ne \beta_j']$ be the Hamming distance between β and β'. The following lemma shows that it is possible to construct a large hypothesis classes where any two models in the hypothesis class are far away under the Hamming distance:

Lemma 6.5 (Raskutti et al. (2011), Lemma 4). *Define $\mathcal{H} = \{z \in \{-1, 0, +1\}^p : \|z\|_0 = s\}$. For p, s even and $s < 2p/3$, there exists a subset $\tilde{\mathcal{H}} \subseteq \mathcal{H}$ with cardinality $|\tilde{\mathcal{H}}| \ge \exp\{\frac{s}{2}\log\frac{p-s}{s/2}\}$ such that $\rho_H(z, z') \ge s/2$ for all dinstinct $s, s' \in \tilde{\mathcal{H}}$.*

Without loss of generality we shall restrain ourselves to even p and $s/2$ scenarios. This does not affect the minimax lower bound to be proved. Using the above lemma and under the condition that $s \le 4p/5$, one can construct Θ consisting of hypothesis of the form in Eq. (13) such that $\log|\Theta| \asymp s\log(p/s)$ and $\|\beta - \beta'\|_2 \ge \sqrt{s/4}\delta$ for all distinct $\beta, \beta' \in \Theta$. It remains to evaluate the KL divergence between P_β and $P_{\beta'}$.

Let x_{obs} and x_{mis} denote the observed and missing covariates of a particular data point and let β_{obs}, β_{mis} be the corresponding partition of coordinates of β. The likelihood of x_{obs} and y can be obtained by integrating out x_{mis} (assuming there are q coordinates that are observed):

$$\begin{aligned}p(y, x_{\text{obs}}; \beta) &= \rho_*^q(1-\rho_*)^{p-q}\int \mathcal{N}_p(x_{\text{obs}}, x_{\text{mis}}; 0, I)\mathcal{N}(y - (x_{\text{obs}}^\top\beta_{\text{obs}} - x_{\text{mis}})^\top\beta_{\text{mis}}; 0, \sigma_\varepsilon^2)\mathrm{d}x_{\text{mis}} \\ &= p(x_{\text{obs}}) \cdot \frac{1}{\sqrt{2\pi(\sigma_\varepsilon^2 + \|\beta_{\text{mis}}\|_2^2)}}\exp\left\{-\frac{(y - x_{\text{obs}}^\top\beta_{\text{obs}})^2}{2(\sigma_\varepsilon^2 + \|\beta_{\text{mis}}\|_2^2)}\right\}.\end{aligned}$$

Note that $p(x_{\rm obs})$ does not depend on β. Subsequently,

$$\begin{aligned}
\mathrm{KL}(P_\beta \| P_{\beta'}) &= \mathbb{E}_{\beta,\rho_*} \log \frac{p(y, x_{\rm obs}; \beta')}{p(y, x_{\rm obs}; \beta)} \\
&= \mathbb{E}_{\beta,\rho_*} \left\{ \frac{1}{2} \log \frac{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2}{\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2} + \frac{1}{2} \left[\frac{(y - x_{\rm obs}^\top \beta'_{\rm obs})^2}{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2} - \frac{(y - x_{\rm obs}^\top \beta_{\rm obs})^2}{\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2} \right] \right\} \\
&= \mathbb{E}_{\rho_*} \left\{ \frac{1}{2} \log \frac{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2}{\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2} + \frac{1}{2} \left[\frac{\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2 + \|\beta_{\rm obs} - \beta'_{\rm obs}\|_2^2}{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2} - 1 \right] \right\} \\
&\overset{(a)}{\leq} \mathbb{E}_{\rho_*} \left\{ \frac{1}{2} \left[\frac{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2}{\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2} + \frac{\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2}{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2} \right] - 1 + \frac{1}{2} \frac{\|\beta_{\rm obs} - \beta'_{\rm obs}\|_2^2}{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2} \right\} \\
&= \mathbb{E}_{\rho_*} \left\{ \frac{1}{2} \frac{\left(\|\beta'_{\rm mis}\|_2^2 - \|\beta_{\rm mis}\|_2^2\right)^2}{(\sigma_\varepsilon^2 + \|\beta_{\rm mis}\|_2^2)(\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2)} + \frac{1}{2} \frac{\|\beta_{\rm obs} - \beta'_{\rm obs}\|_2^2}{\sigma_\varepsilon^2 + \|\beta'_{\rm mis}\|_2^2} \right\}. \qquad (14)
\end{aligned}$$

Here for (a) we apply the inequality that $\log(1+x) \leq x$ for all $x > 0$. For some constant $c \in (0, 1/2)$, define $\mathcal{E}(c)$ as the event that at least $\frac{1-\rho_*}{1+2c}$ portion of the first $s/2$ coordinates in x are missing. By Chernoff bound, [2] $\Pr[\mathcal{E}(c)] \geq 1 - e^{-c^2(1-\rho_*)s}$. Note that under $\mathcal{E}(c)$, $\|\beta_{\rm mis}\|_2^2, \|\beta'_{\rm mis}\|_2^2 \geq \frac{(1-\rho_*)s}{2(1+2c)} a^2$ almost surely. Subsequently,

$$\begin{aligned}
\mathrm{KL}(P_\beta \| P_{\beta'}) &\leq \frac{1}{2} \frac{\mathbb{E}_{\rho_*|\mathcal{E}(c)} \left(\|\beta'_{\rm mis}\|_2^2 - \|\beta_{\rm mis}\|_2^2\right)^2}{\left(\sigma_\varepsilon^2 + \frac{(1-\rho_*)s}{2(1+2c)} a^2\right)^2} + \frac{1}{2} \frac{\mathbb{E}_{\rho_*|\mathcal{E}(c)} \|\beta_{\rm obs} - \beta'_{\rm obs}\|_2^2}{\sigma_\varepsilon^2 + \frac{(1-\rho_*)s}{2(1+2c)} a^2} \\
&+ e^{-c^2(1-\rho_*)s} \left[\frac{1}{2} \frac{\mathbb{E}_{\rho_*|\bar{\mathcal{E}}(c)} \left(\|\beta'_{\rm mis}\|_2^2 - \|\beta_{\rm mis}\|_2^2\right)^2}{\sigma_\varepsilon^4} + \frac{1}{2} \frac{\mathbb{E}_{\rho_*|\bar{\mathcal{E}}(c)} \|\beta_{\rm obs} - \beta'_{\rm obs}\|_2^2}{\sigma_\varepsilon^2} \right].
\end{aligned}$$

Because β and β' are identical in the first $s/2$ coordinates, both $\|\beta'_{\rm mis}\|_2^2 - \|\beta_{\rm mis}\|_2^2$ and $\|\beta_{\rm obs} - \beta'_{\rm obs}\|_2^2$ are independent of $\mathcal{E}(c)$. Therefore,

$$\begin{aligned}
\mathbb{E}_{\rho_*} \left(\|\beta'_{\rm mis}\|_2^2 - \|\beta_{\rm mis}\|_2^2\right)^2 = \mathbb{E}_{\rho_*} \left(\|\beta'_{{\rm mis},>s/2}\|_2^2 - \|\beta_{{\rm mis},>s/2}\|_2^2\right)^2 &\leq 4(1-\rho_*)^2 s^2 \delta^4; \\
\mathbb{E}_{\rho_*} \|\beta'_{\rm obs} - \beta_{\rm obs}\|_2^2 = \mathbb{E}_{\rho_*} \|\beta'_{{\rm obs},>s/2} - \beta_{{\rm obs},>s/2}\|_2^2 &\leq 2\rho_* s \delta^2.
\end{aligned}$$

Here $\beta_{\cdot,>s/2}$ denote the $\beta_\cdot$ vector without its first $s/2$ coordinates, and in both inequalities we note by construction that $\|\beta_{>s/2}\|_0, \|\beta'_{>s/2}\|_0 \leq s/2$. Because $a^2 = \frac{2M^2}{s} - \delta^2$, we have that $\frac{(1-\rho_*)s}{2(1+2c)} a^2 = \frac{1-\rho_*}{1+2c} M^2 - \frac{(1-\rho_*)}{2(1+2c)} s\delta^2$. For now assume that $\frac{1-\rho_*}{1+2c} s\delta^2 \ll \sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c} M^2$, which then implies $\sigma_\varepsilon^2 + \frac{(1-\rho_*)s}{2(1+2c)} a^2 \geq \frac{1}{2} \left(\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c} M^2\right)$. We will justify this assumption at the end of this proof. Combining all inequalities we have

$$\mathrm{KL}(P_\beta \| P_{\beta'}) \leq \frac{8(1-\rho_*)^2 s^2 \delta^4}{\left(\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c} M^2\right)^2} + \frac{2\rho_* s \delta^2}{\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c} M^2} + e^{-c^2(1-\rho_*)s} \left[\frac{2(1-\rho_*)^2 s^2 \delta^4}{\sigma_\varepsilon^4} + \frac{\rho_* s \delta^2}{\sigma_\varepsilon^2} \right].$$

[2] If $X_1, \cdots, X_n$ are i.i.d. random variables taking values in $\{0, 1\}$ then $\Pr[\frac{1}{n} \sum_{i=1}^n X_i < (1-\delta)\mu] \leq \exp\{-\frac{\delta^2 \mu}{2}\}$ for $0 < \delta < 1$, where $\mu = \mathbb{E}X$.

Let P_β^n and $P_{\beta'}^n$ be the distribution of n i.i.d. samples parameterized by β and β', respectively. Because the samples are i.i.d., we have that $\mathrm{KL}(P_\beta^n \| P_{\beta'}^n) = n\mathrm{KL}(P_\beta \| P_{\beta'})$. On the other hand, because $\log|\Theta| \asymp s\log(p/s)$, to ensure $1 - \frac{\mathrm{KL}(P_\beta^n \| P_{\beta'}^n) + \log 1/2}{\log|\Theta|} \geq \Omega(1)$ we only need to show $\mathrm{KL}(P_\beta^n \| P_{\beta'}^n) \asymp s\log(p/s)$, which is implied by

$$
\begin{aligned}
\frac{(1-\rho_*)^2 s^2\delta^4}{\left(\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2\right)^2} \asymp \frac{s\log(p/s)}{n} \quad &\Longleftarrow \quad \delta^2 \asymp \left(\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2\right)\sqrt{\frac{\log(p/s)}{(1-\rho_*)^2 sn}}; \\
\frac{\rho_* s\delta^2}{\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2} \asymp \frac{s\log(p/s)}{n} \quad &\Longleftarrow \quad \delta^2 \asymp \left(\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2\right)\frac{\log(p/s)}{\rho_* n}; \\
e^{-c^2(1-\rho_*)s}\frac{(1-\rho_*)^2 s^2\delta^4}{\sigma_\varepsilon^4} \asymp \frac{s\log(p/s)}{n} \quad &\Longleftarrow \quad \delta^2 \asymp e^{0.5c^2(1-\rho_*)s}\sigma_\varepsilon^2\sqrt{\frac{\log(p/s)}{(1-\rho_*)^2 sn}}; \\
e^{-c^2(1-\rho_*)s}\frac{\rho_* s\delta^2}{\sigma_\varepsilon^2} \quad &\Longleftarrow \quad \delta^2 \asymp e^{c^2(1-\rho_*)s}\sigma_\varepsilon^2\frac{\log(p/s)}{\rho_* n}.
\end{aligned}
$$

Combining all terms we have that

$$
\delta^2 \asymp \min\left\{\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2, e^{0.5c^2(1-\rho_*)s}\sigma_\varepsilon^2\right\} \cdot \min\left\{\sqrt{\frac{\log(p/s)}{(1-\rho_*)^2 sn}}, \frac{\log(p/s)}{\rho_* n}\right\}. \tag{15}
$$

The bound for $\|\beta - \beta'\|_2^2$ can then be obtained by $\|\beta - \beta'\|_2^2 \geq \frac{s}{4}\delta^2$.

The final part of the proof is to justify the assumption that $\frac{1-\rho_*}{1+2c}s\delta^2 \ll \sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2$. Invoking Eq. (15), the assumption is valid if $\frac{1-\rho_*}{1+2c}\max\left\{\sqrt{\frac{s\log(p/s)}{(1-\rho_*)^2 n}}, \frac{s\log(p/s)}{\rho_* n}\right\} \to 0$, which holds if $\frac{s\log(p/s)}{\rho_* n} \to 0$.

6.4 Proof of Theorem 2.3

We again take $\rho_1 = \cdots = \rho_p = \rho_*$. The first term $\frac{\sigma_\varepsilon^2}{\rho_* n}$ in the minimax lower bound is trivial to establish: consider $\beta_0 = \delta e_j$ and $\beta_1 = -\delta e_j$ with $\Sigma_0 = \Sigma_1 = I$. By Eq. (14), we have that

$$
\mathrm{KL}(P_{\beta_0}^n \| P_{\beta_1}^n) = n \cdot \mathrm{KL}(P_{\beta_0} \| P_{\beta_1}) \leq \frac{2\rho_* n\delta^2}{\sigma_\varepsilon^2}.
$$

Equating $\mathrm{KL}(P_{\beta_0}^n \| P_{\beta_1}^n) \asymp O(1)$ we have that $\delta^2 \asymp \frac{\sigma_\varepsilon^2}{\rho_* n}$. Because $\frac{\sigma_\varepsilon^2}{M^2\rho_* n} \to 0$, we know that $\beta_0, \beta_1 \in \mathbb{B}_2(M) \cap \mathbb{B}_0(1)$ when n is sufficiently large. Invoking Le Cam's method (Lemma A.2) with $|\beta_{0j} - \beta_{1j}|^2 = 4\delta^2 \asymp \frac{\sigma_\varepsilon^2}{\rho_* n}$ we prove the desired minimax lower bound of $\frac{\sigma_\varepsilon^2}{\rho_* n}$.

We next focus on the second term in the minimax lower bound that involves $1/\rho_*^2 n$. Without loss of generality assume $j > s-1$. Construct two hypothesis (β_0, Σ_0) and (β_1, Σ_1) as follows:

$$
\beta_0 = (\underbrace{\frac{\tilde{a}}{\sqrt{s-2}}, \cdots, \frac{\tilde{a}}{\sqrt{s-2}}}_{\text{repeat } s-2 \text{ times}}, \tilde{a}, \underbrace{0, \cdots, 0, \tilde{a}\gamma, 0, \cdots, 0}_{\beta_{0j} = \tilde{a}\gamma}), \qquad \Sigma_0 = I_{p\times p} - \gamma(e_{s-1}e_j^\top + e_j e_{s-1}^\top);
$$

$$\beta_1 = (\underbrace{\frac{\tilde{a}}{\sqrt{s-2}}, \cdots, \frac{\tilde{a}}{\sqrt{s-2}}}_{\text{repeat } s-2 \text{ times}}, \tilde{a}, \underbrace{0, \cdots, 0, -\tilde{a}\gamma, 0, \cdots, 0}_{\beta_{0j} = -\tilde{a}\gamma}), \qquad \Sigma_1 = I_{p\times p} + \gamma(e_{s-1}e_j^\top + e_j e_{s-1}^\top).$$

Here $\gamma \to 0$ is some parameter to be determined later and $\tilde{a}$ is set to $\tilde{a} = \sqrt{\frac{M^2}{2+\gamma^2}}$ to ensure that $\|\beta_0\|_2 = \|\beta_1\|_2 = M$. It is immediate by definition that $\beta_0, \beta_1 \in \mathbb{B}_2(M) \cap \mathbb{B}_0(s)$. In addition, by Gershgorin circle theorem all eigenvalues of Σ_0 and Σ_1 lie in $[1-\gamma, 1+\gamma]$. As $\gamma \to 0$, it holds that $\Sigma_0, \Sigma_1 \in \Lambda(\gamma_0)$ for any constant $\gamma_0 \in (0, 1/2)$ when n is sufficiently large. A finite-sample statement of this fact is given at the end of the proof.

Unlike the identity covariance case, the likelihood $p(y, x_{\text{obs}}; \beta, \Sigma)$ for incomplete observations are complicated when Σ has non-zero off-diagonal elements. The following lemma gives a general characterization of the likelihood when $\beta \neq 0$. Its proof is given in the supplementary material.

Lemma 6.6. *Partition the covariance* Σ *as* $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$, *where* Σ_{11} *corresponds to* x_{obs} *and* Σ_{22} *corresponds to* x_{mis}. *Define* $\Sigma_{22:1} = \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}$. *Let* $q = \dim(\Sigma_{11})$ *be the number of observed covariates. Then*

$$\begin{aligned} p(y, x_{\text{obs}}; \beta, \Sigma) \quad = \quad & \rho_*^q (1-\rho_*)^{p-q} \cdot \frac{1}{\sqrt{(2\pi)^q |\Sigma_{11}|}} \exp\left\{ -\frac{1}{2} x_{\text{obs}}^\top \Sigma_{11}^{-1} x_{\text{obs}} \right\} \\ & \cdot \frac{1}{\sqrt{2\pi(\sigma_\varepsilon^2 + \beta_{\text{mis}}^\top \Sigma_{22:1} \beta_{\text{mis}})}} \exp\left\{ -\frac{(y - x_{\text{obs}}^\top \beta_{\text{obs}} - \beta_{\text{mis}}^\top \Sigma_{21}\Sigma_{11}^{-1} x_{\text{obs}})^2}{2(\sigma_\varepsilon^2 + \beta_{\text{mis}}^\top \Sigma_{22:1}\beta_{\text{mis}})} \right\}. \end{aligned}$$

We now present the following lemma, which is key to establish the $1/\rho_*^2$ rate in the minimax lower bound. Its proof is given in the supplementary material.

Lemma 6.7. $p(y, x_{\text{obs}}; \beta_0, \Sigma_0) = p(y, x_{\text{obs}}; \beta_1, \Sigma_1)$ *unless* both x_{s-1} *and* x_j *are observed.*

Let P_0 and P_1 denote the distributions parameterized by (β_0, Σ_0) and (β_1, Σ_1), respectively. Let $\mathcal{A}$ denote the event that both x_{s-1} and x_j are observed. By Lemma 6.7, we have that

$$\text{KL}(P_0 \| P_1) = \Pr[\mathcal{A}] \mathbb{E}_0 \left[\log \frac{p(y, x_{\text{obs}}; \beta_0, \Sigma_0)}{p(y, x_{\text{obs}}; \beta_1, \Sigma_1)} \middle| \mathcal{A} \right] = \rho_*^2 \mathbb{E}_0 \left[\log \frac{p(y, x_{\text{obs}}; \beta_0, \Sigma_0)}{p(y, x_{\text{obs}}; \beta_1, \Sigma_1)} \middle| \mathcal{A} \right].$$

Suppose $\Sigma_0 = [\Sigma_{011} \ \Sigma_{012}; \Sigma_{021} \ \Sigma_{022}]$ and $\Sigma_1 = [\Sigma_{111} \ \Sigma_{112}; \Sigma_{121} \ \Sigma_{122}]$ are partitioned in the same way as in Lemma 6.6. Conditioned on the event $\mathcal{A}$, we have that $\Sigma_{022} = \Sigma_{122} = I_{(p-q)\times(p-q)}$, $\Sigma_{012} = \Sigma_{021}^\top = \Sigma_{112} = \Sigma_{121}^\top = 0_{q\times(p-q)}$, $\Sigma_{011} = I_{q\times q} - \gamma(e_{s-1}e_j^\top + e_j e_{s-1}^\top)$, $\Sigma_{111} = I_{q\times q} + \gamma(e_{s-1}e_j^\top + e_j e_{s-1}^\top)$, and by Lemma A.3, we have that $\Sigma_{011}^{-1} = I + \frac{\gamma^2}{1-\gamma^2}(e_{s-1}e_{s-1}^\top + e_j e_j^\top) + \frac{\gamma}{1-\gamma^2}(e_{s-1}e_j^\top + e_j e_{s-1}^\top)$ and $\Sigma_{111}^{-1} = I + \frac{\gamma^2}{1-\gamma^2}(e_{s-1}e_{s-1}^\top + e_j e_j^\top) - \frac{\gamma}{1-\gamma^2}(e_{s-1}e_j^\top + e_j e_{s-1}^\top)$. In addition, $\det(\Sigma_{011}) = \det(\Sigma_{111}) = 1-\gamma^2$. Note also that $\Sigma_{022:1} = \Sigma_{122:1} = I_{(p-q)\times(p-q)}$ and hence $\beta_{0\text{mis}}^\top \Sigma_{022:1}\beta_{0\text{mis}} = \beta_{1\text{mis}}^\top \Sigma_{122:1}\beta_{1\text{mis}}$ because $\|\beta_{0\text{mis}}\|_2^2 = \|\beta_{1\text{mis}}\|_2^2$ regardless of which covariates are missing. Define $x_{\text{obs},<s} = \{x_j : x_j \text{ is observed}, j < s\}$ and $\beta_{\text{obs},<s} = \{\beta_j : x_j \text{ is observed}, j < s\}$. Subsequently, invoking Lemma 6.6 we get

$$\mathbb{E}_{0|\mathcal{A}} \left[\log \frac{P_0}{P_1} \right] = -\frac{2\gamma}{1-\gamma^2} \mathbb{E}_0[x_{s-1}x_j] - \mathbb{E}_{0|\mathcal{A}} \left\{ \frac{1}{2} \frac{(y - x_{\text{obs}}^\top \beta_{0\text{obs}})^2 - (y - x_{\text{obs}}^\top \beta_{1\text{obs}})^2}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2} \right\}$$

$$
\begin{aligned}
&\overset{(a)}{=} \frac{2\gamma^2}{1-\gamma^2} + \mathbb{E}_{0|\mathcal{A}}\left\{\frac{x_j(\beta_{0j}-\beta_{1j})(y - x_{\text{obs},<s}^\top \beta_{0\text{obs},<s})}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2}\right\} \\
&\overset{(b)}{=} \frac{2\gamma^2}{1-\gamma^2} + \mathbb{E}_{0|\mathcal{A}}\left\{\frac{x_j(\beta_{0j}-\beta_{1j})(x_{\text{mis},<s}^\top \beta_{0\text{mis},<s} + x_j\beta_{0j} + \varepsilon)}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2}\right\} \\
&\overset{(c)}{=} \frac{2\gamma^2}{1-\gamma^2} + \mathbb{E}_{R|\mathcal{A}}\left\{\frac{\beta_{0j}(\beta_{0j}-\beta_{1j})\mathbb{E}_0[x_j^2] + (\beta_{0j}-\beta_{1j})\mathbb{E}_{0|R}[x_j(x_{\text{mis},<s-1}^\top \beta_{0\text{mis},<s-1} + \varepsilon)]}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2}\right\} \\
&= \frac{2\gamma^2}{1-\gamma^2} + \mathbb{E}_{R|\mathcal{A}}\left\{\frac{\beta_{0j}(\beta_{0j}-\beta_{1j})\mathbb{E}_0[x_j^2]}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2}\right\} \\
&= \frac{2\gamma^2}{1-\gamma^2} + \mathbb{E}_{R|\mathcal{A}}\left\{\frac{2\tilde{a}^2\gamma^2}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2}\right\}.
\end{aligned}
$$

Here (a) is due to $\beta_{0\text{obs},<s} = \beta_{1\text{obs},<s}$ and $\beta_{0j}^2 = \beta_{1j}^2$, and (b) is because $\beta_{0k} = 0$ for all $k \geq s$ except for $k = j$. Note also that under $\mathcal{A}$, x_j is observed and hence β_{0j} always belongs to $\beta_{0\text{obs}}$. For (c), note that x_{s-1} is observed under $\mathcal{A}$ and x_j is independent of $x_{<s-1}$ and ε conditioned on R, thanks to the missing completely at random assumption (A3). For any constant $c \in (0, 1/2)$ define $\mathcal{E}'(c)$ as the event that at least $\frac{1-\rho_*}{1+2c}$ portion of the first $(s-2)$ coordinates in x are missing. Note that $\|\beta_{0\text{mis}}\|_2^2 \geq \frac{1-\rho_*}{1+2c}\tilde{a}^2$ almost surely under $\mathcal{A} \cap \mathcal{E}'(C)$ and by Chernoff bound $\Pr[\mathcal{A}] \geq 1 - e^{-c^2(1-\rho_*)(s-2)} \geq 1 - e^{-0.5c^2(1-\rho_*)s}$ for $s \geq 4$. Subsequently, by law of total expectation

$$
\mathbb{E}_{R|\mathcal{A}}\left\{\frac{2\tilde{a}^2\gamma^2}{\sigma_\varepsilon^2 + \|\beta_{0\text{mis}}\|_2^2}\right\} \leq \frac{2\tilde{a}^2\gamma^2}{\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}\tilde{a}^2} + e^{-0.5c^2(1-\rho_*)s}\frac{2\tilde{a}^2\gamma^2}{\sigma_\varepsilon^2}.
$$

Replace $\tilde{a}^2 = \frac{M^2}{2+\gamma^2}$. We then have that

$$
\begin{aligned}
\text{KL}(P_0^n \| P_1^n) &\leq n\rho_*^2\left[\frac{2\gamma^2}{1-\gamma^2} + \frac{2M^2\gamma^2}{(2+\gamma^2)\sigma_\varepsilon^2 + \frac{1-\rho_*}{1+2c}M^2} + e^{-0.5c^2(1-\rho_*)s}\frac{2M^2\gamma^2}{(2+\gamma^2)\sigma_\varepsilon^2}\right] \\
&\leq n\rho_*^2\left[\frac{2\gamma^2}{1-\gamma^2} + \frac{2(1+2c)\gamma^2}{1-\rho_*} + e^{-0.5c^2(1-\rho_*)s}\frac{M^2\gamma^2}{\sigma_\varepsilon^2}\right].
\end{aligned}
$$

Equating $\text{KL}(P_0^n \| P_1^n) \asymp O(1)$ and applying the condition that $\gamma^2 \to 0$, we have that

$$
\gamma^2 \asymp \min\left\{\frac{1-\rho_*}{2(1+2c)}, e^{0.5c^2(1-\rho_*)s}\frac{\sigma_\varepsilon^2}{M^2}\right\}\frac{1}{\rho_*^2 n}. \tag{16}
$$

Subsequently,

$$
\left|\beta_{0j} - \beta_{1j}\right|^2 = 4\tilde{a}^2\gamma^2 \asymp \min\left\{\frac{1-\rho_*}{2(1+2c)}M^2, e^{0.5c^2(1-\rho_*)s}\sigma_\varepsilon^2\right\}\frac{1}{\rho_*^2 n}.
$$

Invoking Lemma A.2 we finish the proof of the minimax lower bound.

Finally, we justify the conditions $\gamma^2 \to 0$ and $\gamma < \gamma_0$ that are used in the proof. Eq. (16) yields $\gamma^2 \leq O(\frac{1}{\rho_*^2 n})$. So $\gamma^2 \to 0$ and $\gamma < \gamma_0$ is implied by $\frac{1}{\gamma_0^2\rho_*^2 n} \to 0$.

6.5 Proof of Theorem 3.1

Using $y = X\beta_0 + \varepsilon$ we have that

$$\tilde{\Sigma}(\hat{\beta}_n - \beta_0) + \left(\frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\hat{\beta}_n\right) = \Big(\underbrace{\frac{1}{n}\tilde{X}^\top X - \tilde{\Sigma}}_{\Delta_n}\Big)\beta_0 + \frac{1}{n}\tilde{X}^\top\varepsilon. \tag{17}$$

Define $\Delta_n = \frac{1}{n}\tilde{X}^\top X - \tilde{\Sigma}$. Recall that $\hat{\beta}_n^u = \hat{\beta}_n + \hat{\Theta}\left(\frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\hat{\beta}_n\right)$. Subsequently, multiplying both sides of Eq. (17) with $\sqrt{n}\hat{\Theta}$ and re-organizing terms we have

$$\begin{aligned}
&\sqrt{n}(\hat{\beta}_n^u - \beta_0) = \sqrt{n}\hat{\Theta}\left(\Delta_n\beta_0 + \frac{1}{n}\tilde{X}^\top\varepsilon\right) - \sqrt{n}(\hat{\Theta}\tilde{\Sigma} - I)(\hat{\beta}_n - \beta_0) \\
&= \sqrt{n}\Sigma_0^{-1}\left(\Delta_n\beta_0 + \frac{1}{n}\tilde{X}^\top\varepsilon\right) - \underbrace{\sqrt{n}(\hat{\Theta}\tilde{\Sigma} - I)(\hat{\beta}_n - \beta_0)}_{r_n} + \underbrace{\sqrt{n}(\hat{\Theta} - \Sigma_0^{-1})\left(\Delta_n\beta_0 + \frac{1}{n}\tilde{X}^\top\varepsilon\right)}_{\tilde{r}_n}.
\end{aligned}$$

Define $r_n = \sqrt{n}(\hat{\Theta}\tilde{\Sigma} - I)(\hat{\beta}_n - \beta_0)$ and $\tilde{r}_n = \sqrt{n}(\hat{\Theta} - \Sigma_0^{-1})\left(\Delta_n\beta_0 + \frac{1}{n}\tilde{X}^\top\varepsilon\right)$

Lemma 6.8. *Suppose* $\frac{\log p}{\rho_*^4 n} \to 0$ *and the conclusion in Lemma 3.1 holds. Then* $\|r_n\|_\infty \le O_{\mathbb{P}}(\sqrt{n}\tilde{\nu}_n\|\hat{\beta}_n - \beta_0\|_1)$ *and* $\|\tilde{r}_n\|_\infty \le O_{\mathbb{P}}(\sigma_x b_0 b_1 \tilde{\nu}_n(\sigma_\varepsilon\sqrt{\frac{\log p}{\rho_*}} + \sigma_x\|\beta_0\|_2\sqrt{\frac{\log p}{\rho_*^2}}))$.

Lemma 6.8 based on Hölder's inequality and is proved in the supplementary materials. If the condition in Eq. (9) holds, Lemma 6.8 implies that $\max\{\|r_n\|_\infty, \|\tilde{r}_n\|_\infty\} \xrightarrow{p} 0$, which means both terms r_n and $\tilde{r}_n$ are asymptotically negligible in the infinity norm sense. It then suffices to analyze the limiting distribution (conditioned on X) of $a_n = \sqrt{n}\Sigma_0^{-1}\left(\Delta_n\beta_0 + \frac{1}{n}\tilde{X}^\top\varepsilon\right)$. By Assumptions (A1) and (A3), $\mathbb{E}\Delta_n|X = 0$, $\mathbb{E}\varepsilon|\tilde{X} = 0$ and hence $\mathbb{E}a_n|X = 0$. We next analyze the conditional covariance $\mathbb{V}a_n|X$. Recall that $\Delta_n = \frac{1}{n}\tilde{X}^\top X - \tilde{\Sigma}$. By definition, for any $j, k \in \{1, \cdots, p\}$

$$[\Delta_n]_{jk} = \begin{cases} \frac{1}{n}\sum_{i=1}^n \frac{R_{ij}}{\rho_j}\left(1 - \frac{R_{ik}}{\rho_k}\right)X_{ij}X_{ik}, & j \neq k; \\ 0, & j = k. \end{cases}$$

Here $R_{ij} = 1$ if X_{ij} is observed and $R_{ij} = 0$ otherwise. Subsequently, $a_n = \Sigma_0^{-1}\tilde{a}_n$ where

$$[\tilde{a}_n]_j = \frac{1}{\sqrt{n}}\sum_{i=1}^n \Bigg(\underbrace{\frac{R_{ij}X_{ij}}{\rho_j}\varepsilon_i + \sum_{k\neq j}\frac{R_{ij}}{\rho_j}\left(1 - \frac{R_{ik}}{\rho_k}\right)X_{ij}X_{ik}\beta_{0k}}_{T_{ij}}\Bigg).$$

Because $R \perp\!\!\!\perp X, \varepsilon$ and $\varepsilon \perp\!\!\!\perp X$, we have that $\mathbb{E}T_{ij}|X = 0$. Therefore, for any $j \in \{1, \cdots, p\}$

$$\mathbb{V}T_{ij}|X = \mathbb{E}\left[|T_{ij}|^2|X\right] = \frac{\sigma_\varepsilon^2 X_{ij}^2}{\rho_j} + \sum_{t\neq j}\frac{1 - \rho_t}{\rho_j\rho_t}X_{ij}^2X_{it}^2\beta_{0t}^2$$

and for $j \neq k$,

$$\mathrm{cov}(T_{ij}, T_{ik}|X) = \mathbb{E}\left[T_{ij}T_{ik}|X\right] = \sigma_\varepsilon^2 X_{ij}X_{ik} + \sum_{t\neq j,k} \frac{1-\rho_t}{\rho_t} X_{ij}X_{ik}X_{it}^2\beta_{0t}^2.$$

Because $\{T_{ij}\}_{i=1}^n$ are i.i.d. random variables, by central limiting theorem, for any subset $S \subseteq [p]$ with constant size

$$[a_n]_{SS} \xrightarrow{d} \mathcal{N}_{|S|}\left(0, \mathrm{cov}_{SS}(a_n|X)\right) \xrightarrow{d} \mathcal{N}_{|S|}\left(0, \left[\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1}\right]_{SS}\right),$$

where all randomness is conditioned on X.

6.6 Proof of Theorem 3.2

By triangle inequality and Hölder's inequality,

$$\begin{aligned}
&\|\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1} - \hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top\|_\infty \\
&\leq \|(\Sigma_0^{-1} - \hat{\Theta})\hat{\Gamma}\Sigma_0^{-1}\|_\infty + \|\hat{\Theta}\hat{\Gamma}(\Sigma_0^{-1} - \hat{\Theta}^\top)\|_\infty + \|\hat{\Theta}(\hat{\Gamma} - \tilde{\Gamma})\hat{\Theta}^\top\|_\infty \\
&\leq 2\max\left\{\|\Sigma_0^{-1}\|_{L_1}, \|\hat{\Theta}\|_{L_1}, \|\hat{\Theta}\|_{L_\infty}\right\} \max\left\{\|\Sigma_0^{-1} - \hat{\Theta}\|_{L_1}, \|\Sigma_0^{-1} - \hat{\Theta}\|_{L_\infty}\right\} \|\hat{\Gamma}\|_\infty + \|\hat{\Theta}\|_{L_1}^2\|\hat{\Gamma} - \tilde{\Gamma}\|_\infty.
\end{aligned}$$

With Lemma 3.1, the bound can be simplified to (with probability $1 - o(1)$)

$$\|\Sigma_0^{-1}\hat{\Gamma}\Sigma_0^{-1} - \hat{\Theta}\tilde{\Gamma}\hat{\Theta}^\top\|_\infty \leq 4b_0b_1^2\tilde{\nu}_n\|\hat{\Gamma}\|_\infty + b_1^2\|\hat{\Gamma} - \tilde{\Gamma}\|_\infty. \tag{18}$$

Note that by standard concentration inequalities of supreme of sub-Gaussian random variables, $\|X\|_\infty \leq O_{\mathbb{P}}(\sigma_x\sqrt{\log p})$. Also, by Hölder's inequality $\|\hat{\Upsilon}\|_\infty \leq \rho_*^{-2}\|X\|_\infty^4\|\beta_0^2\|_1$. Subsequently,

$$\|\hat{\Gamma}\|_\infty \leq \frac{\sigma_\varepsilon^2}{\rho_*}\|X\|_\infty^2 + \frac{\|X\|_\infty^4\|\beta_0\|_2^2}{\rho_*^2} \leq O_{\mathbb{P}}\left\{\sigma_x^4\log^2 p\left(\frac{\sigma_\varepsilon^2}{\sigma_x^2\rho_*} + \frac{\|\beta_0\|_2^2}{\rho_*^2}\right)\right\}. \tag{19}$$

It remains to upper bound $\|\hat{\Gamma} - \tilde{\Gamma}\|_\infty$. Decompose the difference as

$$\|\hat{\Gamma} - \tilde{\Gamma}\|_\infty \leq \sigma_\varepsilon^2\left\|\frac{1}{n}\tilde{X}^\top\tilde{X} - \frac{1}{n}X^\top X - \tilde{D}\mathrm{diag}\left(\frac{1}{n}X^\top X\right)\right\|_\infty + \|\hat{\Upsilon} - \tilde{\Upsilon}\|_\infty.$$

We first focus on the first term. Recall that $\|\tilde{D}\|_\infty \leq 1 - 1/\rho_*$ and $\frac{1}{n}\tilde{X}^\top\tilde{X} = \tilde{\Sigma} + D\mathrm{diag}(\frac{1}{n}\tilde{X}^\top\tilde{X})$. Subsequently, the first infinity norm term is upper bounded by

$$\|\tilde{\Sigma} - \Sigma_0\|_\infty + \left\|D\mathrm{diag}\left(\frac{1}{n}\tilde{X}^\top\tilde{X}\right) - \tilde{D}\mathrm{diag}(\Sigma_0)\right\|_\infty + \frac{1}{\rho_*}\|\hat{\Sigma} - \Sigma_0\|_\infty.$$

By Lemma 6.1, if $\frac{\log p}{\rho_*^4 n} \to 0$ then $\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq O_{\mathbb{P}}(\sigma_x^2\sqrt{\frac{\log p}{\rho_*^2 n}})$ and $\|\hat{\Sigma} - \Sigma_0\|_\infty \leq O_{\mathbb{P}}(\sigma_x^2\sqrt{\frac{\log p}{n}})$. For the remaining term, we invoke the following lemma that is proved in the supplementary materials:

Lemma 6.9. *If* $\frac{\log p}{\rho_* n} \to 0$ *then* $\|D\mathrm{diag}(\frac{1}{n}\tilde{X}^\top\tilde{X}) - \tilde{D}\mathrm{diag}(\Sigma_0)\|_\infty \leq O_{\mathbb{P}}(\sigma_x^2\sqrt{\frac{\log p}{\rho_*^3 n}})$.

Consequently,

$$\sigma_\varepsilon^2 \left\| \frac{1}{n}\tilde{X}^\top \tilde{X} - \frac{1}{n}X^\top X - \tilde{D}\mathrm{diag}\left(\frac{1}{n}X^\top X\right)\right\|_\infty \leq O_{\mathbb{P}}\left\{\sigma_\varepsilon^2\sigma_x^2\sqrt{\frac{\log p}{\rho_*^3 n}}\right\}. \tag{20}$$

Finally, we derive the upper bound for $\|\hat{\Upsilon} - \tilde{\Upsilon}\|_\infty$. We first construct a $p \times p$ matrix $\bar{\Upsilon}$ as an "intermediate" quantity defined as

$$\bar{\Upsilon}_{jk} = \frac{1}{n}\sum_{i=1}^{n}\sum_{t\neq j,k}(1-\rho_t)\tilde{X}_{ij}\tilde{X}_{ik}\tilde{X}_{it}^2\beta_{0t}^2 \quad \text{for} \quad j,k \in \{1,\cdots,p\}.$$

Note that $\bar{\Upsilon}$ involves the missing design $\tilde{X}$ and the true model β_0. Further define $\tilde{\Upsilon}_{jkt}$ and Υ_{jkt} for $j,k,t \in \{1,\cdots,p\}$ as

$$\tilde{\Upsilon}_{jkt} = \frac{1}{n}\sum_{i=1}^{n}(1-\rho_t)\tilde{X}_{ij}\tilde{X}_{ik}\tilde{X}_{it}^2, \qquad \Upsilon_{jkt} = \mathbb{E}\tilde{\Upsilon}_{jkt}|X.$$

We next state the following concentration results on $\tilde{\Upsilon}_{jkt}$ and Υ_{jkt}, which will be proved in the supplementary material.

Lemma 6.10. *Fix $j,k \in [p]$ and suppose $\frac{\log p}{\rho_*^3 n} \to 0$. We then have that*

$$\max_{j,k\in[p]}\max_{t\neq j,k}\left|\Upsilon_{jkt}\right| \leq O_{\mathbb{P}}\left(\frac{\sigma_x^4\log^2 p}{\rho_*^2}\right)$$

and

$$\max_{j,k\in[p]}\max_{t\neq j,k}\left|\tilde{\Upsilon}_{jkt} - \Upsilon_{jkt}\right| \leq O_{\mathbb{P}}\left(\sigma_x^4\log^2 p\sqrt{\frac{\log p}{\rho_*^5 n}}\right).$$

We then upper bound $\|\tilde{\Upsilon} - \hat{\Upsilon}\|_\infty$ by bounding $\|\hat{\Upsilon} - \bar{\Upsilon}\|_\infty$ and $\|\tilde{\Upsilon} - \bar{\Upsilon}\|_\infty$ separately.

Upper bound for $\|\tilde{\Upsilon} - \bar{\Upsilon}\|_\infty$ By definition, $\tilde{\Upsilon}_{jk} = \sum_{t\neq j,k}\tilde{\Upsilon}_{jkt}\hat{\beta}_{nt}^2$ and $\bar{\Upsilon}_{jkt} = \sum_{t\neq j,k}\tilde{\Upsilon}_{jkt}\beta_{0t}^2$. Hölder's inequality then yields

$$\|\tilde{\Upsilon} - \bar{\Upsilon}\|_\infty \leq \max_{j,k\in[p]}\max_{t\neq j,k}\left|\tilde{\Upsilon}_{jkt}\right| \cdot \|\hat{\beta}_n^2 - \beta_0^2\|_1.$$

Under the condition that $\frac{\log p}{\rho_*^3 n} \to 0$, it holds that $\max_{j,k}\max_{t\neq j,k}|\tilde{\Upsilon}_{jkt}| \leq O_{\mathbb{P}}(1)\cdot\max_{j,k}\max_{t\neq j,k}|\Upsilon_{jkt}|$. Furthermore, $\|\hat{\beta}_n^2 - \beta_0^2\|_1 \leq \|\hat{\beta}_n + \beta_0\|_\infty\|\hat{\beta}_n - \beta_0\|_1 \leq (\|\beta_0\|_2 + \|\hat{\beta}_n - \beta_0\|_2)\|\hat{\beta}_n - \beta_0\|_1$. Invoking Lemma 6.10 and the condition that $\|\hat{\beta}_n - \beta_0\|_2 \xrightarrow{p} 0$ we get

$$\|\tilde{\Upsilon} - \bar{\Upsilon}\|_\infty \leq O_{\mathbb{P}}\left\{\frac{\sigma_x^4\log^2 p}{\rho_*^2}\|\beta_0\|_2\|\hat{\beta}_n - \beta_0\|_1\right\}. \tag{21}$$

Upper bound for $\|\hat{\Upsilon} - \bar{\Upsilon}\|_\infty$ Note that $\hat{\Upsilon}_{jk} = \sum_{t \neq j,k} \Upsilon_{jkt}\beta_{0t}^2$ and $\bar{\Upsilon}_{jk} = \sum_{t \neq j,k} \tilde{\Upsilon}_{jkt}\beta_{0t}^2$. By Hölder's inequality,

$$\|\bar{\Upsilon} - \hat{\Upsilon}\|_\infty \leq \max_{j,k \in [p]} \max_{t \neq j,k} \left|\tilde{\Upsilon}_{jkt} - \Upsilon_{jkt}\right| \cdot \|\beta_0^2\|_1.$$

Invoking Lemma 6.10 we then have

$$\|\bar{\Upsilon} - \hat{\Upsilon}\|_\infty \leq O_{\mathbb{P}}\left\{\sigma_x^4 \log^2 p \|\beta_0\|_2^2 \sqrt{\frac{\log p}{\rho_*^5 n}}\right\}. \tag{22}$$

Finally, combining Eqs. (18,19,20,21,22) we complete the proof of Theorem 3.2.

Appendix A Technical Lemmas

Lemma A.1 (generalized Fano's inequality, (Ibragimov & Has' minskii, 2013))**.** *Let Θ be a parameter set and $d : \Theta \times \Theta \to \mathbb{R}_{\geq 0}$ be a semimetric. Let P_θ be the distribution induced by θ and P_θ^n be the distribution of n i.i.d. observations from P_θ. If $d(\theta, \theta') \geq \alpha$ and $\mathrm{KL}(P_\theta \| P_{\theta'}) \leq \beta$ for all distinct $\theta, \theta' \in \Theta$, then*

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta} \mathbb{E}_{P_{\theta^n}}\left[d(\hat{\theta}, \theta)\right] \geq \frac{\alpha}{2}\left(1 - \frac{n\beta + \log 2}{\log |\Theta|}\right).$$

Lemma A.2 (Le Cam's method, (Le Cam, 2012))**.** *Suppose P_{θ_0} and P_{θ_1} are distributions induced by θ_0 and θ_1. Let $P_{\theta_0}^n$ and $P_{\theta_1}^n$ be distributions of n i.i.d. observations from P_{θ_0} and P_{θ_1}, respectively. Then for any estimator $\hat{\theta}$ it holds that*

$$\frac{1}{2}\left[\Pr_{P_{\theta_0}^n}(\hat{\theta} \neq \theta_0) + \Pr_{P_{\theta_1}^n}(\hat{\theta} \neq \theta_1)\right] \geq \frac{1}{2} - \frac{1}{2}\|P_{\theta_0}^n - P_{\theta_1}^n\|_{\mathrm{TV}} \geq \frac{1}{2} - \frac{1}{2\sqrt{2}}\sqrt{n\mathrm{KL}(P_{\theta_0} \| P_{\theta_1})}.$$

Lemma A.3 (Miller (1981), Eq. (13))**.** *Suppose H is a matrix of rank at most 2 and $(I + H)$ is invertible. Then*

$$(I + H)^{-1} = I - \frac{aH - H^2}{a + b},$$

where $a = 1 + \mathrm{tr}(H)$ and $2b = [\mathrm{tr}(H)]^2 + \mathrm{tr}(H^2)$.

References

Bach, F. R. (2008). Consistency of the group lasso and multiple kernel learning. *Journal of Machine Learning Research*, *9*(Jun), 1179–1225.

Belloni, A., Rosenbaum, M., & Tsybakov, A. B. (2016). Linear and conic programming estimators in high dimensional errors-in-variables models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.

Bickel, P. J., Ritov, Y., & Tsybakov, A. B. (2009). Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, (pp. 1705–1732).

Cai, T., Liu, W., & Luo, X. (2011). A constrained L1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, *106*(494), 594–607.

Cai, T. T., & Guo, Z. (2015). Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. *arXiv preprint arXiv:1506.05539*.

Cai, T. T., Liang, T., & Rakhlin, A. (2014). Geometric inference for general high-dimensional linear inverse problems. *arXiv preprint arXiv:1404.4408*.

Candes, E., & Tao, T. (2007). The dantzig selector: Statistical estimation when p is much larger than n. *The Annals of Statistics*, (pp. 2313–2351).

Candès, E. J., Romberg, J., & Tao, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, *52*(2), 489–509.

Chen, Y., & Caramanis, C. (2013). Noisy and missing data regression: Distribution-oblivious support recovery. In *International Conference on Machine Learning*.

Datta, A., & Zou, H. (2015). Cocolasso for high-dimensional error-in-variables regression. *arXiv preprint arXiv:1510.07123*.

Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on information theory*, *52*(4), 1289–1306.

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *The Annals of statistics*, *32*(2), 407–499.

Fan, J., & Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, *96*(456), 1348–1360.

Ibragimov, I. A., & Has' minskii, R. Z. (2013). *Statistical estimation: asymptotic theory*, vol. 16. Springer Science & Business Media.

Javanmard, A., & Montanari, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *Journal of Machine Learning Research*, *15*(1), 2869–2909.

Le Cam, L. (2012). *Asymptotic methods in statistical decision theory*. Springer Science & Business Media.

Loh, P.-L., & Wainwright, M. (2012a). High-dimensional regression with noisy and missing data: provable guarantees with nonconvexity. *The Annals of Statistics*, *40*(3), 1637–1664.

Loh, P.-L., & Wainwright, M. J. (2012b). Corrupted and missing predictors: Minimax bounds for high-dimensional linear regression. In *2012 IEEE International Symposium on Information Theory Proceedings (ISIT)*.

Loh, P.-L., & Wainwright, M. J. (2015). Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *Journal of Machine Learning Research*, *16*, 559–616.

Miller, K. S. (1981). On the inverse of the sum of matrices. *Mathematics Magazine*, *54*(2), 67–72.

Nielsen, T. O., West, R. B., Linn, S. C., Alter, O., Knowling, M. A., O'Connell, J. X., Zhu, S., Fero, M., Sherlock, G., Pollack, J. R., et al. (2002). Molecular characterisation of soft tissue tumours: a gene expression study. *The Lancet*, *359*(9314), 1301–1307.

Raskutti, G., Wainwright, M. J., & Yu, B. (2011). Minimax rates of estimation for high-dimensional linear regression over-balls. *IEEE Transactions on Information Theory*, *57*(10), 6976–6994.

Rosenbaum, M., & Tsybakov, A. (2010). Sparse recovery under matrix uncertainty. *The Annals of Statistics*, *38*(5), 2620–2651.

Rosenbaum, M., & Tsybakov, A. (2013). Improved matrix uncertainty selector.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, (pp. 267–288).

van de Geer, S., Bühlmann, P., Ritov, Y., & Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*, *32*(3), 1166–1202.

Wainwright, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1-constrained quadratic programming (lasso). *IEEE transactions on information theory*, *55*(5), 2183–2202.

Zhang, C.-H., & Zhang, S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, *76*, 217–242.

Zhao, P., & Yu, B. (2006). On model selection consistency of lasso. *Journal of Machine learning research*, *7*(Nov), 2541–2563.

Supplementary Material for: Rate Optimal Estimation and Confidence Intervals for High-dimensional Regression with Missing Covariates

Yining Wang[1], Jialei Wang[3], Sivaraman Balakrishnan[1,2], and Aarti Singh[1]

[1]Machine Learning Department, Carnegie Mellon University
[2]Department of Statistics, Carnegie Mellon University
[3]Department of Computer Science, University of Chicago

February 10, 2017

This supplementary material provides detailed proofs for technical lemmas whose proofs are omitted in the main text.

A Proofs of concentration bounds

A.1 Proof of Lemma 6.1

Fix arbitrary $u, v \in \mathcal{S}$. For $j, k \in [p]$ and $\ell \in \{0, 1, 2\}$, define

$$\xi_{jk}^{(0)}(R_i, \rho) = 1, \quad \xi_{jk}^{(1)}(R_i, \rho) = \frac{R_{ij}}{\rho_j}, \quad \xi_{jk}^{(2)}(R_i, \rho) = \begin{cases} \frac{R_{ij}}{\rho_j}, & j = k; \\ \frac{R_{ij}R_{ik}}{\rho_j \rho_k}, & j \neq k. \end{cases}$$

Also let $T_i^{(\ell)} = \sum_{j,k=1}^p \xi_{jk}(R_i, \rho) X_{ij} X_{ik} u_j v_k$. We then have that

$$\left|u^\top (\hat{\Sigma} - \Sigma_0) v\right| = \left|\frac{1}{n}\sum_{i=1}^n T_i^{(0)} - \mathbb{E}T_i^{(0)}\right|, \tag{S1}$$

$$\left|u^\top (\frac{1}{n}\tilde{X}^\top X - \Sigma_0) v\right| = \left|\frac{1}{n}\sum_{i=1}^n T_i^{(1)} - \mathbb{E}T_i^{(2)}\right|, \tag{S2}$$

$$\left|u^\top (\tilde{\Sigma} - \Sigma_0) v\right| = \left|\frac{1}{n}\sum_{i=1}^n T_i^{(2)} - \mathbb{E}T_i^{(2)}\right|. \tag{S3}$$

The main idea is to use Berstein inequality with moment conditions (Lemma E.5) to establish concentration bounds and achieve optimal dependency over ρ. Define $V^{(\ell)} = \mathbb{E}\left[|T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}|^2\right]$. We then have that

$$V^{(\ell)} \leq \mathbb{E}|T_i^{(\ell)}|^2 = \sum_{j,k,j',k'=1}^p \mathbb{E}\left\{\xi_{jk}^{(\ell)}\xi_{j'k'}^{(\ell)}\right\} \mathbb{E}\left\{X_{ij}X_{ik}X_{ij'}X_{ik'}u_j v_k u_{j'} v_{k'}\right\}.$$

It is then of essential importance to evaluate $\mathbb{E}\left\{\xi_{jk}^{(\ell)}\xi_{j'k'}^{(\ell)}\right\}$. For $\ell = 0$ the expectation trivially equals 1. For $\ell = 1$ and $\ell = 2$, we apply the following proposition, which is easily proved by definition.

Proposition A.1. $\mathbb{E}\left\{\xi_{jk}^{(1)}\xi_{j'k'}^{(1)}\right\} = 1 + I[j = j'](\frac{1}{\rho_j} - 1)$ *and* $\mathbb{E}\left\{\xi_{jk}^{(2)}\xi_{j'k'}^{(2)}\right\} = 1 + I[j = j'](\frac{1}{\rho_j} - 1) + I[k = k'](\frac{1}{\rho_k} - 1) + I[j = j' \wedge k = k'](\frac{1}{\rho_j} - 1)(\frac{1}{\rho_k} - 1)$. *Here $I[\cdot]$ is the indicator function.*

We are now ready to derive $\mathbb{E}|T_i^{(\ell)}|^2$.

$$
\begin{aligned}
\mathbb{E}|T_i^{(0)}|^2 &= \mathbb{E}\left\{|X_i^\top u|^2|X_i^\top v|^2\right\};\\
\mathbb{E}|T_i^{(1)}|^2 &= \mathbb{E}\left\{|X_i^\top u|^2|X_i^\top v|^2\right\} + \sum_{j=1}^p \left(\frac{1}{\rho_j} - 1\right) u_j^2 \mathbb{E}\left\{X_{ij}^2|X_i^\top v|^2\right\}\\
&\le \mathbb{E}\left\{|X_i^\top u|^2|X_i^\top v|^2\right\} + \frac{1}{\rho_*}\sum_{j=1}^p u_j^2 \mathbb{E}\left\{|X_i^\top e_j|^2|X_i^\top v|^2\right\};\\
\mathbb{E}|T_i^{(2)}|^2 &= \mathbb{E}\left\{|X_i^\top u|^2|X_i^\top v|^2\right\} + \sum_{j=1}^p \left(\frac{1}{\rho_j} - 1\right)(u_j^2 + v_j^2)\mathbb{E}\left\{X_{ij}^2|X_i^\top v|^2\right\}\\
&\quad + \sum_{j,k=1}^p \left(\frac{1}{\rho_j} - 1\right)\left(\frac{1}{\rho_k} - 1\right) u_j^2 v_k^2 \mathbb{E}\left\{X_{ij}^2 X_{ik}^2\right\}\\
&\le \mathbb{E}\left\{|X_i^\top u|^2|X_i^\top v|^2\right\} + \frac{1}{\rho_*}\sum_{j=1}^p (u_j^2 + v_j^2)\mathbb{E}\left\{|X_i^\top e_j|^2|X_i^\top v|^2\right\} + \frac{1}{\rho_*^2}\sum_{j,k=1}^p u_j^2 v_k^2 \mathbb{E}\left\{|X_i^\top e_j|^2|X_i^\top e_k|^2\right\}.
\end{aligned}
$$

By Cauchy-Schwarts inequality and moment upper bounds of sub-Gaussian random variables (Lemma E.1), we have that

$$
\mathbb{E}\left\{|X_i^\top a|^2|X_i^\top b|^2\right\} \le \sqrt{\mathbb{E}|X_i^\top a|^4}\sqrt{\mathbb{E}|X_i^\top b|^4} \le 16\sigma_x^4\|a\|_2^2\|b\|_2^2.
$$

Consequently, there exists universal constant $c_2 > 0$ such that

$$
\mathbb{E}|T_i^{(0)}|^2 \le c_2\sigma_x^4\|u\|_2^2\|v\|_2^2, \quad \mathbb{E}|T_i^{(1)}|^2 \le \frac{c_2}{\rho_*}\sigma_x^4\|u\|_2^2\|v\|_2^2, \quad \mathbb{E}|T_i^{(2)}|^2 \le \frac{c_2}{\rho_*^2}\sigma_x^4\|u\|_2^2\|v\|_2^2.
$$

We next establish $L > 0$ so that the moment condition in Lemma E.5 is satisfied, namely $\mathbb{E}|T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}|^k \le \frac{1}{2}V^{(\ell)}L^{k-2}k!$ for all $k > 1$. Note that for all $\ell \in \{0, 1, 2\}$, there exist functions $\xi_j^{(\ell)}$ and $\bar{\xi}_j^{(\ell)}$ *only depending on* j such that $\xi_{jk}^{(\ell)} = \xi_j^{(\ell)}\xi_k^{(\ell)} + I[j = k]\cdot\bar{\xi}_j^{(\ell)}$ and furthermore $\max_j|\xi_j^{(\ell)}| \le 1/\rho_*$, $\bar{\xi}_j^{(0)} = \bar{\xi}_j^{(1)} = 0$ and $\max_j|\bar{\xi}_j^{(2)}| \le 1/\rho_*^2$. Subsequently,

$$
\begin{aligned}
\mathbb{E}|T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}|^k &= \mathbb{E}\left|\sum_{j,k=1}^p \left(\xi_j^{(\ell)}\xi_k^{(\ell)} + I[j = k]\cdot\bar{\xi}_j^{(\ell)} - 1\right) X_{ij}X_{ik}u_jv_k\right|^k\\
&\le 3^k\left(\mathbb{E}\left|\sum_{j,k=1}^p \xi_j^{(\ell)}\xi_k^{(\ell)}X_{ij}X_{ik}u_jv_k\right|^k + \mathbb{E}\left|\sum_{j=1}^p \bar{\xi}_j^{(\ell)}X_{ij}^2u_jv_j\right|^k + \mathbb{E}\left|\sum_{j,k=1}^p X_{ij}X_{ik}u_jv_k\right|^k\right).
\end{aligned}
$$

Here the second line is a consequence of the following inequality: for all $a, b, c \geq 0$ we have that $(a+b+c)^k \leq (3\max\{a,b,c\})^k \leq 3^k \max\{a^k, b^k, v^k\} \leq 3^k(a^k+b^k+c^k)$. Define $\tilde{u}_j = u_j \xi_j^{(\ell)}$, $\tilde{v}_k = v_k \xi_k^{(\ell)}$, $\bar{u}_j = u_j \sqrt{|\bar{\xi}_j^{(\ell)}|}$ and $\bar{v}_j = v_j \sqrt{|\bar{\xi}_j^{(\ell)}|}$. Apply Lemma E.6 with $|\sum_{j=1}^p \bar{\xi}_j^{(\ell)} X_{ij}^2 u_j v_j| \leq X_i^\top A X_i$, $A = \mathrm{diag}(|\bar{u}_1 \bar{v}_1|, \cdots, |\bar{u}_p \bar{v}_p|)$ and note that $\mathrm{tr}(A) = \sqrt{\mathrm{tr}(A^2)} = |\bar{u}|^\top |\bar{v}| \leq \|\bar{u}\|_2 \|\bar{v}\|_2$ and $\|A\|_{\mathrm{op}} = \max_{1 \leq j \leq p} |\bar{u}_j \bar{v}_j| \leq \|\bar{u}\|_2 \|\bar{v}\|_2$. Subsequently, for all $t > 0$

$$\Pr\left[X_i^\top A X_i > 3\sigma_x^2 \|\bar{u}\|_2 \|\bar{v}\|_2 (1+t)\right] \leq e^{-t}. \tag{S4}$$

Let $F(x) = \Pr[X_i^\top A X_i \leq x], x \geq 0$ be the CDF of $X_i^\top A X_i$ and $G(x) = 1 - F(x)$. Using integration by parts, we have that

$$\mathbb{E}|X_i^\top A X_i|^k = \int_0^\infty x^k \mathrm{d}F(x) = -\int_0^\infty x^k \mathrm{d}G(x) = \int_0^\infty k x^{k-1} G(x) \mathrm{d}x.$$

Here in the last equality we use the fact that $\lim_{x \to \infty} x^k G(x) = 0$ for any fixed $k \in \mathbb{N}$, because $G(x) \leq \exp\{1 - \frac{x}{M}\}$ by Eq. (S4), where $M = 3\sigma_x^2 \|\bar{u}\|_2 \|\bar{v}\|_2$. Consequently,

$$\begin{aligned}
\mathbb{E}|X_i^\top A X_i|^k &= \int_0^M k x^{k-1} G(x) \mathrm{d}x + k \int_M^\infty x^{k-1} G(x) \mathrm{d}x \\
&\leq M^k + k \int_0^\infty M^{k-1} (1+z)^{k-1} e^{-z} \cdot M \mathrm{d}z \\
&= M^k + k M^k \int_0^\infty (1+z)^{k-1} e^{-z} \mathrm{d}z \\
&\leq M^k + k M^k \cdot k! \leq (k+1)! M^k.
\end{aligned}$$

Here in the second line we apply change-of-variable $x = M(1+z)$ and the fact that $G(M(1+z)) \leq e^{-z}$ in the integration term. Because $2^k \geq k+1$ for all $k \geq 1$, we conclude that

$$\mathbb{E}\left|\sum_{j=1}^p \bar{\xi}_j^{(\ell)} X_{ij}^2 u_j v_j\right|^k \leq 6^k \sigma_x^{2k} k! \mathbb{E}\|\bar{u}\|_2^k \|\bar{v}\|_2^k, \quad \forall k \geq 1.$$

Subsequently, applying Cauchy-Schwarts inequality together with moment bounds for sub-Gaussian random variables (Lemma E.1) we obtain

$$\begin{aligned}
&\mathbb{E}|T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}|^k \\
&\leq 3^k \left(\sqrt{\mathbb{E}|X^\top \tilde{u}|^{2k}} \sqrt{\mathbb{E}|X^\top \tilde{v}|^{2k}} + 6^k \sigma_x^{2k} k! \mathbb{E}\|\bar{u}\|_2^k \|\bar{v}\|_2^k + \sqrt{\mathbb{E}|X^\top u|^{2k}} \sqrt{\mathbb{E}|X^\top v|^{2k}}\right) \\
&\leq 3^k \cdot 2k \cdot 6^k \Gamma(k) \sigma_x^{2k} \cdot \left(\sqrt{\mathbb{E}\|\tilde{u}\|_2^{2k}} \sqrt{\mathbb{E}\|\tilde{v}\|_2^{2k}} + \sqrt{\mathbb{E}\|\bar{u}\|_2^{2k}} \sqrt{\mathbb{E}\|\bar{v}\|_2^{2k}} + \|u\|_2^k \|v\|_2^k\right) \\
&\leq \rho_*^{\ell/2} \left(\frac{C' \|u\|_2 \|v\|_2 \sigma_x^2}{\rho_*^\ell}\right)^k k!,
\end{aligned}$$

where $C' < \infty$ is some absolute constant. Compare the bound of $\mathbb{E}|T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}|^k$ with the variance $\mathbb{E}|T_i^{(\ell)}|^2$ we obtained earlier, we have that $L = \sigma_x^2 \|u\|_2 \|v\|_2 \cdot C'^3 / \rho_*^{1.5\ell}$ is sufficient to guarantee

$\mathbb{E}|T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}|^k \leq \frac{1}{2}V^{(\ell)}L^{k-2}k!$ for all [1] $k > 2$. Applying Bernstein inequality with moment conditions (Lemma E.5) and union bound over all $u, v \in \mathcal{S}$, we have that

$$\Pr\left[\forall u, v \in \mathcal{S}, \left|\frac{1}{n}\sum_{i=1}^{n} T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}\right| > \|u\|_2\|v\|_2\epsilon\right] \leq 2N^2 \exp\left\{-\frac{n\epsilon^2}{2(\tilde{V}^{(\ell)} + \tilde{L}\epsilon)}\right\}$$

for all $\epsilon > 0$, where $\tilde{V}^{(\ell)} = \frac{V^{(\ell)}}{\|u\|_2^2\|v\|_2^2}$ and $\tilde{L} = \frac{L}{\|u\|_2\|v\|_2}$. Subsequently,

$$\begin{aligned}\sup_{u,v\in\mathcal{S}}\left|\frac{1}{n}\sum_{i=1}^{n} T_i^{(\ell)} - \mathbb{E}T_i^{(\ell)}\right| &\leq O_{\mathbb{P}}\left(\|u\|_2\|v\|_2 \max\left\{\frac{\tilde{L}\log N}{n}, \sqrt{\frac{\tilde{V}^{(\ell)}\log N}{n}}\right\}\right) \\ &\leq O_{\mathbb{P}}\left(\sigma_x^2\|u\|_2\|v\|_2 \max\left\{\frac{\log N}{\rho_*^{1.5\ell}n}, \sqrt{\frac{\log N}{\rho_*^{\ell}n}}\right\}\right).\end{aligned}$$

A.2 Proof of Lemma 6.2

Define $\delta_j = \frac{1}{n}\sum_{i=1}^{n} Z_{ij}$ where $Z_{ij} = \widetilde{X}_{ij}\varepsilon_i$. Because $\mathbb{E}\varepsilon_i|X = 0$, we have that $\mathbb{E}Z_{ij} = 0$. In addition,

$$\mathbb{E}|Z_{ij}|^2 = \frac{\sigma_\varepsilon^2\sigma_x^2}{\rho_j} \leq \frac{\sigma_\varepsilon^2\sigma_x^2}{\rho_*} =: V$$

and

$$\begin{aligned}\mathbb{E}|Z_{ij}|^k &= \rho_j \cdot \frac{1}{\rho_j^k} \cdot \mathbb{E}\varepsilon_i^k \cdot \mathbb{E}|X_{ij}|^k \\ &\leq \frac{1}{\rho_*^{k-1}} \cdot k^2 2^k \sigma_x^k \sigma_\varepsilon^k \Gamma\left(\frac{k}{2}\right)^2 \\ &\leq \frac{k^2(2\sigma_x\sigma_\varepsilon)^k}{\rho_*^{k-1}} k! \\ &\leq \rho_* \left(\frac{4\sigma_x\sigma_\varepsilon}{\rho_*}\right)^k k!.\end{aligned}$$

By setting $L = 64\sigma_x\sigma_\varepsilon/\rho_*$ we have that $\mathbb{E}|Z_{ij}|^k \leq \frac{1}{2}VL^{k-2}k!$ for all $k > 1$. Subsequently, applying Bernstein inequality with moment conditions (Lemma E.5) and union bound over $j = 1, \cdots, p$ we have that

$$\Pr\left[\|\delta\|_\infty > \epsilon\right] \leq 2p \exp\left\{-\frac{n\epsilon^2}{2(V + L\epsilon)}\right\}$$

for any $\epsilon > 0$. Suppose $\frac{\epsilon L}{V} \to 0$. We then have that

$$\|\delta\|_\infty \leq O_{\mathbb{P}}\left(\sigma_\varepsilon\sigma_x\sqrt{\frac{\log p}{\rho_* n}}\right).$$

The condition $\frac{\epsilon L}{V} \to 0$ is satisfied with $\frac{\log p}{\rho_* n} \to 0$.

[1] The case of $k = 2$ is trivially true.

A.3 Proof of Lemma 6.9

Fix arbitrary $j \in \{1, \cdots, p\}$ and consider

$$T_{ij} = (1-\rho_j)\tilde{X}_{ij}^2 = \frac{(1-\rho_j)R_{ij}X_{ij}^2}{\rho_j^2}.$$

It is easy to verify that $[D\mathrm{diag}(\frac{1}{n}\tilde{X}^\top\tilde{X})]_{jj} = \frac{1}{n}\sum_{i=1}^n T_{ij}$ and $[\tilde{D}\mathrm{diag}(\Sigma_0)]_{jj} = \frac{1}{n}\sum_{i=1}^n \mathbb{E}T_{ij} = \frac{(1-\rho_j)\Sigma_{0jj}}{\rho_j}$. We use moment based Bernstein's inequality (Lemma E.5) to bound the perturbation $|\frac{1}{n}\sum_{i=1}^n T_{ij} - \mathbb{E}T_{ij}|$. Define $V_j = \mathbb{E}|T_{ij} - \mathbb{E}T_{ij}|^2$. We then have

$$V_j \le \mathbb{E}|T_{ij}|^2 = \frac{(1-\rho_j)^2\mathbb{E}X_{ij}^4}{\rho_j^3} \le \frac{3\sigma_x^4}{\rho_*^3}$$

and for all $k \ge 3$,

$$\mathbb{E}|T_{ij} - \mathbb{E}T_{ij}|^k \le 2^k\left(\mathbb{E}|T_{ij}|^k + |\mathbb{E}T_{ij}|^k\right) \le \frac{4^{k+1}}{\rho_*^{2k-1}}\sigma_x^{2k}k!.$$

It can then be verified that $\mathbb{E}|T_{ij} - \mathbb{E}T_{ij}|^k \le \frac{1}{2}V_jL^{k-2}k!$ for all $k \ge 2$ if $L = \frac{512\sigma_x^2}{\rho_*^2}$. By Lemma E.5 and a union bound over all $j \in \{1, \cdots, p\}$, we have that

$$\Pr\left[\forall j, \left|\frac{1}{n}\sum_{i=1}^n T_{ij} - \mathbb{E}T_{ij}\right| > \epsilon\right] \le 2p\exp\left\{-\frac{n\epsilon^2}{2(V+L\epsilon)}\right\}$$

for all $\epsilon > 0$, where $V = \frac{3\sigma_x^4}{\rho_*^3}$ and $L = \frac{512\sigma_x^2}{\rho_*^2}$. Under the assumption that $\frac{\epsilon L}{V} \to 0$, we have that

$$\left\|\tilde{D}\mathrm{diag}\left(\frac{1}{n}\tilde{X}^\top\tilde{X}\right) - D\mathrm{diag}(\Sigma_0)\right\|_\infty = \sup_{1\le j\le p}\left|\frac{1}{n}\sum_{i=1}^n T_{ij} - \mathbb{E}T_{ij}\right| \le O_{\mathbb{P}}\left(\sigma_x^2\sqrt{\frac{\log p}{\rho_*^3 n}}\right).$$

The condition $\frac{\epsilon L}{V} \to 0$ is then satisfied with $\frac{\log p}{\rho_* n} \to 0$.

A.4 Proof of Lemma 6.10

By definition and the missing data model,

$$\Upsilon_{jkt} = \mathbb{E}\tilde{\Upsilon}_{jkt}|X = \begin{cases} \frac{1}{n}\sum_{i=1}^n \frac{1-\rho_t}{\rho_j\rho_k}X_{ij}^2X_{it}^2, & j = k; \\ \frac{1}{n}\sum_{i=1}^n \frac{1-\rho_t}{\rho_k}X_{ij}X_{ik}X_{it}^2, & j \ne k. \end{cases}$$

Subsequently,

$$\max_{j,k\in[p]}\max_{t\ne j,k}\left|\Upsilon_{jkt}\right| \le \frac{\|X\|_\infty^4}{\rho_*^2} \le O_{\mathbb{P}}\left(\frac{\sigma_x^4\log^2 p}{\rho_*^2}\right).$$

To prove the second part of this lemma, we first fix arbitrary $j, k \in [p]$ and $t \ne j, k$. Define

$$T_{ijkt} = \left(\xi_{jkt}(R_i, \rho) - \mathbb{E}\xi_{jkt}(R_i, \rho)\right)X_{ij}X_{ik}X_{it}^2,$$

where $\xi_{jkt}(R_i, \rho) = \frac{(1-\rho_t)R_{ij}R_{ik}R_{it}}{\rho_j\rho_k\rho_t^2}$. It is easy to verify that $\tilde{\Upsilon}_{jkt} - \Upsilon_{jkt} = \frac{1}{n}\sum_{i=1}^n T_{ijkt}$ and $\mathbb{E}T_{ijkt}|X = 0$. We then use Bernstein inequality with support conditions (Lemma E.4) to bound the concentration of $\frac{1}{n}\sum_{i=1}^n T_{ijkt}$ towards zero. Define $A = \max_{i,j,k,t}|T_{ijkt}|$ and $V = \max_{i,j,k,t}\mathbb{E}|T_{ijkt}|^2$. By Hölder's inequality we have that

$$A \le \frac{\|X\|_\infty^4}{\rho_*^4} \le O_{\mathbb{P}}\left(\frac{\sigma_x^4 \log^2 p}{\rho_*^4}\right).$$

Here in the $O_{\mathbb{P}}(\cdot)$ notation the randomness is on the generating process of X and is *independent* of the randomness of missing patterns R. In addition, note that

$$\mathbb{E}\left|\left(\xi_{jkt} - \mathbb{E}\xi_{jkt}\right)\left(\xi_{jkt'} - \mathbb{E}\xi_{jkt'}\right)\right| \le \frac{1}{\rho_*^5}$$

for all $j, k, t, t' \in \{1, \cdots, p\}$ and $t, t' \neq j, k$. Subsequently,

$$V = \max_{i,j,k,t}\mathbb{E}|T_{ijkt}|^2 \le \frac{1}{\rho_*^5}X_{ij}^2X_{ik}^2X_{it}^4 \le \frac{\|X\|_\infty^8}{\rho_*^5} \le O_{\mathbb{P}}\left(\frac{\sigma_x^8\log^4 p}{\rho_*^5}\right).$$

Applying Lemma E.4 conditioned on $\|X\|_\infty \le O(\frac{\sigma_x^4\log^2 p}{\rho_*^2})$, we have that with probability $1 - O(\delta)$ for some $\delta = o(1)$ the following holds:

$$\left|\frac{1}{n}\sum_{i=1}^n T_{ijkt}\right| \le O\left(\sigma_x^4\log^2 p\sqrt{\frac{\log(1/\delta)}{\rho_*^5 n}}\right) =: \epsilon,$$

provided that $\frac{\epsilon A}{V} \to 0$. Applying union bound over all $j, k \in [p]$ and $t \in [p]\backslash\{j,k\}$ we get

$$\max_{j,k\in[p]}\max_{t\neq j,k}\left|\frac{1}{n}\sum_{i=1}^n T_{ijkt}\right| \le O_{\mathbb{P}}\left(\sigma_x^4\log^2 p\sqrt{\frac{\log p}{\rho_*^5 n}}\right),$$

The condition $\frac{\epsilon A}{V} \to 0$ is satisfied with $\frac{\log p}{\rho_*^3 n} \to 0$.

B Proof of restricted eigenvalue conditions

Lemma B.1. *Suppose A, B are $p \times p$ random matrices with $\Pr[\|A - B\|_\infty \le M] \ge 1 - o(1)$ for some $M < \infty$. If A satisfies $\mathrm{RE}(s, \phi_{\min})$ and B satisfies $\mathrm{RE}(s, \phi'_{\min})$, then with probability $1 - o(1)$ we have that*

$$\phi'_{\min} \ge \phi_{\min} - \{O(1)\cdot\varphi_{u,v}(A, B; O(s\log(Mp))) + O(1/n)\}.$$

Proof. For any $h \in \mathbb{R}^p$ it holds that

$$\frac{h^\top Bh}{h^\top h} \ge \frac{h^\top Ah}{h^\top h} - \frac{h^\top(B-A)h}{h^\top h}.$$

With appropriate scalings, it suffices to bound

$$\sup_{h:\|h_{J^c}\|_1\le\|h_J\|_1,\|h\|_2\le 1}\left|h^\top(B-A)h\right|$$

for all $J\subseteq[p]$, $|J|\le s$ as the largest possible gap between $\phi_{\min}$ and $\phi'_{\min}$.

Define $\mathbb{B}_p(r)=\{x\in\mathbb{R}^p:\|x\|_p\le r\}$ as the p-norm ball of radius r. Because $\|h_{J^c}\|_1\le\|h_J\|_1$ implies $\|h\|_1\le 2\|h_J\|_1\le 2\sqrt{s}\|h\|_2$, we have that

$$\sup_{h:\|h_{J^c}\|_1\le\|h_J\|_1,\|h\|_2\le 1}\left|h^\top(B-A)h\right|\le\sup_{h\in\mathbb{B}_2(1)\cap\mathbb{B}_1(2\sqrt{s})}\left|h^\top(B-A)h\right|.$$

By Lemma 11 in the supplementary material of Loh & Wainwright (2012a), we have that

$$\begin{aligned}\mathbb{B}_2(1)\cap\mathbb{B}_1(2\sqrt{s}) &\subseteq 3\mathrm{conv}\left\{\mathbb{B}_0(4s)\cap\mathbb{B}_2(1)\right\}\\ &\subseteq \mathrm{conv}\{\underbrace{\mathbb{B}_0(4s)\cap\mathbb{B}_2(3)}_{K(4s)}\}.\end{aligned}$$

Here $\mathrm{conv}(A)$ denotes the convex hull of set A. Let $K(4s)=\mathbb{B}_0(4s)\cap\mathbb{B}_2(3)$ and denote $N_{\epsilon,\|\cdot\|_2}(K(4s))$ as the *covering number* of $K(4s)$ with respect to the Euclidean norm $\|\cdot\|_2$. That is, $N_{\epsilon,\|\cdot\|_2}(K(4s))$ is the size of the smallest *covering set* $H\subseteq K(4s)$ such that $\sup_{h\in K(4s)}\inf_{h'\in H}\|h-h'\|_2\le\epsilon$. By definition of the concentration bounds, we have that with probability $1-o(1)$

$$\sup_{h\in H}\left|h^\top(A-B)h\right|\le\varphi_{u,u}(A,B;\log|H|)\sup_{h\in H}\|h\|_2^2\le 9\varphi_{u,u}(A,B;\log N_{\epsilon,\|\cdot\|_2}(K(4s))).$$

Subsequently, for any $\epsilon\in(0,1)$ with probability $1-o(1)$

$$\begin{aligned}\sup_{h\in\mathbb{B}_2(1)\cap\mathbb{B}_1(2\sqrt{s})}\left|h^\top(B-A)h\right| &\le \sup_{h\in\mathrm{conv}\{K(4s)\}}\left|h^\top(A-B)h\right|\\ &\le \sup_{\substack{\xi_1,\cdots,\xi_T\ge 0,\\ \xi_1+\cdots+\xi_T=1,\\ h_1,\cdots,h_T\in K(4s)}}\sum_{i,j=1}^{T}\xi_i\xi_j\left|h_i^\top(A-B)h_j\right|\\ &\le \sup_{h,h'\in K(4s)}\left|h^\top(A-B)h'\right|\\ &\le \sup_{h,h'\in H_{\epsilon,\|\cdot\|_2}[K(4s)]}\left|h^\top(A-B)h'\right|+(6\epsilon+3\epsilon^2)\|A-B\|_{L_2}\\ &\le 36\left\{\varphi_{u,u}(A,B;\log N_{\epsilon,\|\cdot\|_2}(K(4s)))+\epsilon pM\right\}.\end{aligned}$$

Here the last inequality is implied by the condition that $\|A-B\|_\infty\le M$ with probability $1-O(n^{-\alpha})$. Taking $\epsilon=O(1/(p^2M))$ we have that $\epsilon pM=O(1/p)=O(1/n)$.

The final part of the proof is to establish upper bounds for the covering number $N_{\epsilon,\|\cdot\|_2}(K(4s))$. First note that by definition

$$K(4s)=\bigcup_{J\subseteq[p]:|J|\le 4s}\left\{h:\mathrm{supp}(h)=J\wedge\|h\|_2\le 3\right\}.$$

The covering number of a union of subsets can be upper bounded by the following proposition:

Proposition B.1. *Let $K = K_1 \cup \cdots \cup K_m$. Then $N_{\epsilon,\|\cdot\|_2}(K) \leq \sum_{i=1}^m N_{\epsilon,\|\cdot\|_2}(K_i)$.*

Proof. Let $H_i \subseteq K_i$ be covering sets of subset K_i. Define $H = H_1 \cup \cdots \cup H_m$. Clearly $|H| \leq \sum_{i=1}^m |H_i| \leq \sum_{i=1}^m N_{\epsilon,\|\cdot\|_2}(K_i)$. It remains to prove that H is a valid ϵ-covering set of K. Take arbitrary $h \in K$. By definition, there exists $i \in [m]$ such that $h \in K_i$. Subsequently, there exists $h^* \in H_i \subseteq H$ such that $\|h - h^*\|_2 \leq \epsilon$. Therefore, H is a valid ϵ-covering set of K. □

Define $K_J(r) = \{h : \mathrm{supp}(h) = J \wedge \|h\|_2 \leq r\}$. The covering number of K_J is established in the following proposition:

Proposition B.2. $N_{\epsilon,\|\cdot\|_2}(K_J(r)) \leq \left(\frac{4r+\epsilon}{\epsilon}\right)^{|J|}$.

Proof. $K_J(r)$ is nothing but a centered $|J|$-dimensional ball of radius r, locating at the coordinates indexed by J. The covering number result of high-dimensional ball is due to Lemma 2.5 of van de Geer (2010). □

Combining the three propositions, we obtain

$$\log N_{\epsilon,\|\cdot\|_2}(K(4s)) \leq \log\left(\sum_{j=0}^{4s}\binom{p}{j}\right) + \log\left\{\left(\frac{12+\epsilon/2}{\epsilon/2}\right)^{4s}\right\} \leq O\left(s\log(p/\epsilon)\right).$$

With the configuration of $\epsilon = O(1/(p^2M))$, we have that

$$\log N_{\epsilon,\|\cdot\|_2}(K(4s)) \leq O(s\log(pM)).$$

□

We are now ready to prove Lemma 6.4.

Proof of Lemma 6.4. Consider $A = \tilde{\Sigma}$ and $B = \Sigma_0$ in Lemma B.1. Lemma 6.1 yields

$$\varphi_{u,v}(\tilde{\Sigma}, \Sigma_0; O(s\log(Mp))) \leq O\left(\sigma_x^2 \max\left\{\frac{s\log(Mp)}{\rho_*^3 n}, \sqrt{\frac{s\log(Mp)}{\rho_*^2 n}}\right\}\right) =: \epsilon.$$

By Lemma B.1, to prove this corollary it is sufficient to show that $\frac{\epsilon}{\lambda_{\min}(\Sigma_0)} \to 0$. Note also that $M = \|\tilde{\Sigma} - \Sigma_0\|_\infty \leq \frac{\|X\|_\infty}{\rho_*^2} \leq O\left(\frac{\sigma_x\sqrt{\log p}}{\rho_*^2}\right)$ with probability $1 - o(1)$. The condition $\frac{\epsilon}{\lambda_{\min}(\Sigma_0)} \to 0$ can then be satisfied with $\frac{\sigma_x^4 s\log(\sigma_x \log p/\rho_*)}{\rho_*^3\lambda_{\min}^2 n} \to 0$. □

C Proof of Lemma 3.1

Lemma C.1. *Suppose $\frac{\log p}{\rho_*^4 n} \to 0$ and $\tilde{\nu}_n \asymp \sigma_x^2 b_1\sqrt{\frac{\log p}{\rho_*^2 n}}$. Then with probability $1 - o(1)$ the population precision matrix Σ_0^{-1} is a feasible solution to Eq. (7); that is, $\max\{\|\tilde{\Sigma}\Sigma_0^{-1} - I_{p\times p}\|_\infty, \|\Sigma_0^{-1}\tilde{\Sigma} - I_{p\times p}\|_\infty\} \leq \nu_n$.*

Proof. First by Hölder's inequality we have that

$$\|\tilde{\Sigma}\Sigma_0^{-1} - I\|_\infty = \|(\tilde{\Sigma} - \Sigma_0)\Sigma_0^{-1}\|_\infty \leq \|\Sigma_0^{-1}\|_{L_1}\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq b_1\|\tilde{\Sigma} - \Sigma_0\|_\infty.$$

By Lemma 6.1, with probability $1 - o(1)$

$$\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq \varphi_{u,v}\left(\tilde{\Sigma}, \Sigma_0; 2\log p\right) \leq O\left(\sigma_x^2\sqrt{\frac{\log p}{\rho_*^2 n}}\right),$$

provided that $\frac{\log p}{\rho_*^4 n} \to 0$. Subsequently, we have that

$$\|\Sigma_0^{-1}\|_{L_1}\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq O\left(\sigma_x^2 b_1\sqrt{\frac{\log p}{\rho_*^2 n}}\right) \leq \tilde{\nu}_n \tag{S5}$$

with probability $1 - o(1)$. The $\|\Sigma_0^{-1}\tilde{\Sigma} - I\|_\infty$ term can be bounded in the same way by noting that $\|\Sigma_0^{-1}\tilde{\Sigma} - I\|_\infty \leq \|\Sigma_0^{-1}\|_{L_\infty}\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq b_1\|\tilde{\Sigma} - \Sigma_0\|_\infty$. □

Lemma C.2. *Suppose Σ_0^{-1} is a feasible solution to the CLIME optimization problem in Eq. (7). Then $\max\{\|\hat{\Theta}\|_{L_1}, \|\hat{\Theta}\|_{L_\infty}\} \leq \|\Sigma_0^{-1}\|_{L_1}$ and $\|\hat{\Theta} - \Sigma_0^{-1}\|_\infty \leq 2\tilde{\nu}_n\|\Sigma_0^{-1}\|_{L_1}$.*

Proof. We first establish that $\|\hat{\Theta}\|_{L_1} \leq \|\Sigma_0^{-1}\|_{L_1}$. In Cai et al. (2011) it is proved that the solution set of Eq. (7) is identical to the solution set of

$$\hat{\Theta} = \{\hat{\omega}_i\}_{i=1}^p, \quad \hat{\omega}_i \in \operatorname{argmin}_{\omega_i \in \mathbb{R}^p}\left\{\|\omega_i\|_1 : \|\tilde{\Sigma}\omega_i - e_i\|_\infty \leq \tilde{\nu}_n\right\}.$$

Because Σ_0^{-1} belongs to the feasible set of the above constrained optimization problem, we have that $\|\hat{\omega}_i\|_1 \leq \|\Sigma_0^{-1}\|_{L_1}$ for all $i = 1, \cdots, p$ and hence $\|\hat{\Theta}\|_{L_1} \leq \|\Sigma_0^{-1}\|_{L_1}$. The inequality $\|\hat{\Theta}\|_{L_\infty} \leq \|\Sigma_0^{-1}\|_{L_1}$ can be proved by applying the same argument to $\hat{\Theta}^\top$.

We next prove the infinity norm bound fot the estimation error $\hat{\Theta} - \Sigma_0^{-1}$. By triangle inequality,

$$\|\Sigma_0(\hat{\Theta} - \Sigma_0^{-1})\|_\infty \leq \|\tilde{\Sigma}\hat{\Theta} - I\|_\infty + \|(\tilde{\Sigma} - \Sigma_0)\hat{\Theta}\|_\infty \leq \tilde{\nu}_n + \|(\tilde{\Sigma} - \Sigma_0)\hat{\Theta}\|_\infty.$$

Using Hölder's inequality, we have that

$$\|(\tilde{\Sigma} - \Sigma_0)\hat{\Theta}\|_\infty \leq \|\hat{\Theta}\|_{L_1}\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq \|\Sigma_0^{-1}\|_{L_1}\|\tilde{\Sigma} - \Sigma_0\|_\infty \leq \tilde{\nu}_n.$$

Here the last inequality is due to Eq. (S5). Subsequently, $\|\Sigma_0(\hat{\Theta} - \Sigma_0^{-1})\|_\infty \leq 2\tilde{\nu}_n$. Applying Hölder's inequality again we obtain

$$\|\hat{\Theta} - \Sigma_0^{-1}\|_\infty \leq \|\Sigma_0^{-1}\|_{L_1}\|(\tilde{\Sigma} - \Sigma_0)\hat{\Theta}\|_\infty \leq 2\tilde{\nu}_n\|\Sigma_0^{-1}\|_{L_1}.$$

□

To translate the infinity-norm estimation error Σ_0^{-1} into an L_1-norm bound that we desire, we need the following lemma that establishes basic inequality of the estimation error:

Lemma C.3. *Suppose Σ_0^{-1} is a feasible solution to Eq. (7). Then under Assumption (A5) we have that $\max\{\|\hat{\Theta} - \Sigma_0^{-1}\|_{L_1}, \|\hat{\Theta} - \Sigma_0^{-1}\|_{L_\infty}\} \leq 2b_0\|\hat{\Theta} - \Sigma_0^{-1}\|_\infty$.*

Proof. Let $\hat{\omega}_i$ and $\hat{\omega}_{0i}$ be the ith columns of $\hat{\Theta}$ and Σ_0^{-1}, respectively. Let J_i denote the support size of $\hat{\omega}_{0i}$. Definte $\hat{h} = \hat{\omega}_i - \omega_{0i}$. We then have that

$$\|\hat{\omega}_i\|_1 = \|\omega_{0i} + \hat{h}_{J_i^c}\|_1 + \|\hat{h}_{J_i}\|_1 \geq \|\omega_{0i}\|_1 - \|\hat{h}_{J_i^c}\|_1 + \|\hat{h}_{J_i}\|_1.$$

On the other hand, $\|\hat{\omega}_i\|_1 \leq \|\omega_{0i}\|_1$ as shown in the proof of Lemma C.2. Subsequently, $\|\hat{h}_{J_i^c}\|_1 \leq \|\hat{h}_{J_i}\|_1$ and hence

$$\|\hat{\omega}_i - \omega_{0i}\|_1 = 2\|\hat{h}_{J_i}\|_1 \leq 2|J_i|\|\hat{h}\|_\infty \leq 2b_0\|\hat{\omega}_i - \omega_{0i}\|_\infty.$$

Because the above inequality holds for all $i = 1, \cdots, p$, we conclude that $\|\hat{\Theta} - \Sigma_0^{-1}\|_{L_1} \leq 2b_0\|\hat{\Theta} - \Sigma_0^{-1}\|_\infty$. The bound for $\|\hat{\Theta} - \Sigma_0^{-1}\|_{L_\infty}$ can be proved by applying the same argument to $\hat{\Theta}^\top$. □

Combining all the above lemmas, we have that with probability $1 - o(1)$

$$\max\{\|\hat{\Theta} - \Sigma_0^{-1}\|_{L_1}, \|\hat{\Theta} - \Sigma_0^{-1}\|_{L_\infty}\} \leq 2\tilde{\nu}_n b_0 \|\Sigma_0^{-1}\|_{L_1} \leq O\left\{\sigma_x^2 b_0 b_1^2 \sqrt{\frac{\log p}{\rho_*^2 n}}\right\}.$$

D Proofs of the other technical lemmas

D.1 Proof of Lemma 6.3

We first show that under the conditions on n, $\tilde{\lambda}_n$ and $\check{\lambda}_n$ specified in the lemma, the true regression vector β_0 is feasible to both optimization problems with high probability; that is, $\|\frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\beta_0\|_\infty \leq \tilde{\lambda}_n$ and $\|\frac{1}{n}\tilde{X}^\top y - \Sigma_0\beta_0\|_\infty \leq \check{\lambda}_n$ with probability $1 - o(1)$.

Consider $\hat{\beta}_n$ first. Apply $y = X\beta_0 + \varepsilon$ and Definition 6.1, we have that with probability $1 - o(1)$

$$\begin{aligned}\left\|\frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\beta_0\right\|_\infty &\leq \left\|\left(\frac{1}{n}\tilde{X}^\top X - \Sigma_0\right)\beta_0\right\|_\infty + \left\|\left(\tilde{\Sigma} - \Sigma_0\right)\beta_0\right\|_\infty + \left\|\frac{1}{n}\tilde{X}^\top\varepsilon\right\|_\infty \\ &\leq \left\{\varphi_{u,v}\left(\frac{1}{n}\tilde{X}^\top X, \Sigma_0; \log p\right) + \varphi_{u,v}\left(\tilde{\Sigma}, \Sigma_0; \log p\right)\right\}\|\beta_0\|_2 + \sigma_\varepsilon^2\varphi_{\varepsilon,\infty}\left(\frac{1}{n}\tilde{X}\right).\end{aligned}$$

Now apply Lemmas 6.1 and 6.2: with probability $1 - o(1)$

$$\left\|\frac{1}{n}\tilde{X}^\top y - \tilde{\Sigma}\beta_0\right\|_\infty \leq O\left\{\sigma_x\sqrt{\frac{\log p}{n}}\left(\frac{\sigma_x\|\beta_0\|_2}{\rho_*} + \frac{\sigma_\varepsilon}{\sqrt{\rho_*}}\right)\right\} \leq \tilde{\lambda}_n,$$

provided that $\frac{\log p}{\rho_*^4 n} \to 0$. The same line of argument applies to the second inequality by the following decomposition: under the condition that $\frac{\log p}{\rho_*^2 n} \to 0$, with probability $1 - o(1)$

$$\begin{aligned}\left\|\frac{1}{n}\tilde{X}^\top y - \Sigma_0\beta_0\right\|_\infty &\leq \left\|\left(\frac{1}{n}\tilde{X}^\top X - \Sigma_0\right)\beta_0\right\|_\infty + \left\|\frac{1}{n}\tilde{X}^\top\varepsilon\right\|_\infty \\ &\leq \varphi_{u,v}\left(\frac{1}{n}\tilde{X}^\top X, \Sigma_0, \log p\right)\|\beta_0\|_2 + \sigma_\varepsilon^2\varphi_{\varepsilon,\infty}\left(\frac{1}{n}\tilde{X}\right) \\ &\leq O\left\{\sigma_x\sqrt{\frac{\log p}{\rho_* n}}\left(\sigma_x\|\beta_0\|_2 + \sigma_\varepsilon\right)\right\} \leq \check{\lambda}_n.\end{aligned}$$

We are now ready to prove Lemma 6.3. We only prove the assertion involving $\hat{\beta}_n$, because the same argument applies for $\check{\beta}_n$ as well. Let $\hat{h} = \hat{\beta}_n - \beta_0$. Because $J_0 = \mathrm{supp}(\beta_0)$, we have that

$$\|\hat{\beta}_n\|_1 = \|\beta_0 + \hat{h}_{J_0}\|_1 + \|\hat{h}_{J_0^c}\|_1 \geq \|\beta_0\|_1 - \|\hat{h}_{J_0}\|_1 + \|\hat{h}_{J_0^c}\|_1.$$

On the other hand, because both $\hat{\beta}_n$ and β_0 are feasible, by definition of the optimization problem we have that $\|\hat{\beta}_n\|_1 \leq \|\beta_0\|_1$. Combining both chains of inequalities we arrive at $\|\hat{h}_{J_0^c}\|_1 \leq \|\hat{h}_{J_0}\|_1$, which is to be demonstrated.

D.2 Proof of Lemma 6.6

Proposition D.1. *Suppose $X \sim \mathcal{N}(\mu, \nu^2)$ for $\mu \in \mathbb{R}$ and $\nu > 0$. Then for any $b \in \mathbb{R}$ and $a > 0$, it holds that*

$$\mathbb{E}\frac{1}{\sqrt{2\pi a^2}} \exp\left\{-\frac{(X-b)^2}{2a^2}\right\} = \sqrt{\frac{\nu^2}{a^2+\nu^2}} \exp\left\{-\frac{(\mu-b)^2}{2(a^2+\nu^2)}\right\}.$$

Proof. Because $X \sim \mathcal{N}(\mu, \nu^2)$,

$$\begin{aligned}
&\sqrt{2\pi\nu^2}\mathbb{E} \exp\left\{-\frac{(X-b)^2}{2a^2}\right\} \\
&= \int \exp\left\{-\frac{(x-\mu)^2}{2\nu^2} - \frac{(x-b)^2}{2a^2}\right\} \mathrm{d}x \\
&= \int \exp\left\{-\frac{(a^2+\nu^2)x^2 - 2(a^2\mu+\nu^2 b)x + a^2\mu^2 + \nu^2 b^2}{2a^2\nu^2}\right\} \mathrm{d}x \\
&= \int \exp\left\{-\frac{1}{2a^2\nu^2}\left[(a^2+\nu^2)\left(x - \frac{a^2\mu+\nu^2 b}{a^2+\nu^2}\right)^2 - \frac{(a^2\mu+\nu^2 b)^2}{a^2+\nu^2} + \nu^2 b^2 + a^2\mu^2\right]\right\} \mathrm{d}x \\
&= \exp\left\{-\frac{(\mu-b)^2}{2(a^2+\nu^2)}\right\} \int \exp\left\{-\frac{a^2+\nu^2}{2a^2\nu^2}\left(x - \frac{a^2\mu+\nu^2 b}{a^2+\nu^2}\right)^2\right\} \mathrm{d}x \\
&= \exp\left\{-\frac{(\mu-b)^2}{2(a^2+\nu^2)}\right\} \sqrt{\frac{2\pi a^2\nu^2}{a^2+\nu^2}}.
\end{aligned}$$

The proposition is then proved by multiplying both sides by $\sqrt{2\pi a^2/\nu^2}$. □

We now consider the likelihood $p(y, x_{\mathrm{obs}}; \beta, \Sigma)$. Integrating out the missing parts x_{mis} we have

$$\begin{aligned}
p(y, x_{\mathrm{obs}}; \beta, \Sigma) &= p(x_{\mathrm{obs}}) \int \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} \exp\left\{-\frac{(y - x_{\mathrm{obs}}^\top \beta_{\mathrm{obs}} - x_{\mathrm{mis}}^\top \beta_{\mathrm{mis}})^2}{2\sigma\varepsilon^2}\right\} \mathrm{d}P(x_{\mathrm{mis}}|x_{\mathrm{obs}}) \\
&= p(x_{\mathrm{obs}})\mathbb{E}_u\left[\exp\left\{-\frac{(y - x_{\mathrm{obs}}^\top \beta_{\mathrm{obs}} - u)^2}{2\sigma_\varepsilon^2}\right\} \middle| x_{\mathrm{obs}}\right],
\end{aligned}$$

where $u = x_{\mathrm{mis}}^\top \beta_{\mathrm{mis}}$ follows conditional distribution $u|x_{\mathrm{obs}} \sim \mathcal{N}(\mu, \nu^2)$ with $\mu = x_{\mathrm{obs}}^\top \Sigma_{12}\Sigma_{22}^{-1}\beta_{\mathrm{mis}}$ and $\nu^2 = \beta_{\mathrm{mis}}^\top \Sigma_{22:1}\beta_{\mathrm{mis}}$. Applying Proposition D.1 with $a = \sigma$ and $b = y - x_{\mathrm{obs}}^\top \beta_{\mathrm{obs}}$, we have

$$\mathbb{E}_u\left[\exp\left\{-\frac{(y-x_{\text{obs}}^\top\beta_{\text{obs}}-u)^2}{2\sigma_\varepsilon^2}\right\}\middle| x_{\text{obs}}\right]$$
$$=\frac{1}{\sqrt{2\pi(\sigma_\varepsilon^2+\beta_{\text{mis}}^\top\Sigma_{22:1}\beta_{\text{mis}})}}\exp\left\{-\frac{(y-x_{\text{obs}}^\top\beta_{\text{obs}}-\beta_{\text{mis}}^\top\Sigma_{21}\Sigma_{11}^{-1}x_{\text{obs}})^2}{2(\sigma_\varepsilon^2+\beta_{\text{mis}}^\top\Sigma_{22:1}\beta_{\text{mis}})}\right\}.$$

Finally, $R \perp\!\!\!\perp x$, $x_{\text{obs}} \sim \mathcal{N}_q(0,\Sigma_{11})$ and hence

$$p(x_{\text{obs}})=\rho^q(1-\rho)^{p-q}\cdot\frac{1}{\sqrt{(2\pi)^q|\Sigma_{11}|}}\exp\left\{-\frac{1}{2}x_{\text{obs}}^\top\Sigma_{11}^{-1}x_{\text{obs}}\right\}.$$

D.3 Proof of Lemma 6.7

We prove this lemma by discussing three cases separately when at least one covariate of x_{s-1} and x_j are missing. Assume in each case Σ_0 and Σ_1 are partitioned as in Lemma 6.6; that is, $\Sigma_0=[\Sigma_{011}\ \ \Sigma_{012};\Sigma_{021}\ \ \Sigma_{022}]$ and $\Sigma_1=[\Sigma_{111}\ \ \Sigma_{112};\Sigma_{121}\ \ \Sigma_{122}]$.

1. *Both x_{s-1} and x_j are missing.* In this case $\Sigma_{011}=\Sigma_{111}=I_{q\times q}$ and $\Sigma_{012}=\Sigma_{112}=\Sigma_{021}^\top=\Sigma_{121}^\top=0_{q\times(p-q)}$. Therefore, $\Sigma_{011}=\Sigma_{111}$ and the first two terms in $p(y,x_{\text{obs}};\beta_0,\Sigma_0)$ and $p(y,x_{\text{obs}};\beta_1,\Sigma_1)$ are identical. In addition, $\Sigma_{022:1}=\Sigma_{022}=I-\gamma(e_{s-1}e_j^\top+e_je_{s-1}^\top)$ and $\Sigma_{122:1}=\Sigma_{122}=I+\gamma(e_{s-1}e_j^\top+e_je_{s-1}^\top)$. Subsequently, $\beta_{0\text{mis}}^\top\Sigma_{022:1}\beta_{0\text{mis}}=\|\beta_{0\text{mis}}\|_2^2-2\gamma\beta_{0,s-1}\beta_{0j}=\|\beta_{0\text{mis}}\|_2^2-2\tilde{a}^2\gamma^2$, $\beta_{1\text{mis}}^\top\Sigma_{122:1}\beta_{1\text{mis}}=\|\beta_{1\text{mis}}\|_2^2+2\gamma\beta_{1,s-1}\beta_{1j}=\|\beta_{1\text{mis}}\|_2^2-2\tilde{a}^2\gamma^2$. Because $\|\beta_{0\text{mis}}\|_2^2=\|\beta_{1\text{mis}}\|_2^2$ regardless of which covariates are missing, we have that $\beta_{0\text{mis}}^\top\Sigma_{022:1}\beta_{0\text{mis}}=\beta_{1\text{mis}}^\top\Sigma_{122:1}\beta_{1\text{mis}}$ and hence the last term in $p(y,x_{\text{obs}};\beta_0,\Sigma_0)$ and $p(y,x_{\text{obs}};\beta_1,\Sigma_1)$ are identical, because $\beta_{0\text{mis}}^\top\Sigma_{021}\Sigma_{011}^{-1}=\beta_{1\text{mis}}^\top\Sigma_{121}\Sigma_{111}^{-1}=0$ and $\beta_{0\text{obs}}=\beta_{1\text{obs}}$ when x_j is missing.

2. *x_{s-1} is observed but x_j is missing.* In this case, $\Sigma_{011}=\Sigma_{111}=I_{q\times q}$, $\Sigma_{022}=\Sigma_{122}=I_{(p-q)\times(p-q)}$, $\Sigma_{012}=\Sigma_{021}^\top=-\gamma e_{s-1}e_j^\top$ and $\Sigma_{112}=\Sigma_{121}^\top=\gamma e_{s-1}e_j^\top$. Therefore, $\Sigma_{011}=\Sigma_{111}=I$ and hence the first two terms in the likelihood are identical. In addition, $\Sigma_{022:1}=I-\gamma^2e_je_j^\top=\Sigma_{122:1}$ and hence $\beta_{0\text{mis}}^\top\Sigma_{022:1}\beta_{0\text{mis}}=\beta_{1\text{mis}}^\top\Sigma_{122:1}\beta_{1\text{mis}}=\|\beta_{\text{mis}}\|_2^2-\tilde{a}^2\gamma^4$. Finally, $\beta_{0\text{obs}}=\beta_{1\text{obs}}$ when x_j is missing and $\beta_{0\text{mis}}^\top\Sigma_{021}\Sigma_{011}^{-1}=\beta_{1\text{mis}}^\top\Sigma_{121}\Sigma_{111}^{-1}=-\tilde{a}^2\gamma^2$. Therefore the last term in both likelihoods are the same as well.

3. *x_j is observed but x_{s-1} is missing.* In this case, $\Sigma_{011}=\Sigma_{111}=I_{q\times q}$, $\Sigma_{022}=\Sigma_{122}=I_{(p-q)\times(p-q)}$, $\Sigma_{012}=\Sigma_{021}^\top=-\gamma e_je_{s-1}^\top$ and $\Sigma_{112}=\Sigma_{121}^\top=\gamma e_je_{s-1}^\top$. Therefore, $\Sigma_{011}=\Sigma_{111}=I$ and hence the first two terms in the likelihood are identical. In addition, $\Sigma_{022:1}=I-\gamma^2e_{s-1}e_{s-1}^\top=\Sigma_{122:1}$ and hence $\beta_{0\text{mis}}^\top\Sigma_{022:1}\beta_{0\text{mis}}=\beta_{1\text{mis}}^\top\Sigma_{122:1}\beta_{1\text{mis}}=\|\beta_{\text{mis}}\|_2^2-\tilde{a}^2\gamma^2$. Finally, $\beta_{0\text{obs}}^\top x_{\text{obs}}+\beta_{0\text{mis}}^\top\Sigma_{021}\Sigma_{011}^{-1}x_{\text{obs}}=\beta_{0\text{obs},<s}^\top x_{\text{obs},<s}+\beta_{0j}x_j-\gamma\beta_{0,s-1}x_j=\beta_{0\text{obs},<s}^\top x_{\text{obs},<s}$ because $\beta_{0j}=\tilde{a}\gamma$ and $\beta_{0,s-1}=\tilde{a}$. Similarly, $\beta_{1\text{obs}}^\top x_{\text{obs}}+\beta_{1\text{mis}}^\top\Sigma_{121}\Sigma_{111}^{-1}x_{\text{obs}}=\beta_{1\text{obs},<s}^\top x_{\text{obs},<s}+\beta_{1j}x_j+\gamma\beta_{1,s-1}x_j=\beta_{1\text{obs},<s}^\top x_{\text{obs},<s}$. Because $\beta_{0\text{obs},<s}=\beta_{1\text{obs},<s}$, we conclude that the last term of both likelihoods are the same.

D.4 Proof of Lemma 6.8

We first prove the upper bound for $\|r_n\|_\infty$. By Hölder's inequality,

$$\|r_n\|_\infty\le\sqrt{n}\|\hat{\Theta}\tilde{\Sigma}-I\|_\infty\|\hat{\beta}_n-\beta_0\|_1\le\sqrt{n}\tilde{\nu}_n\|\hat{\beta}_n-\beta_0\|_1,$$

where the last inequality is due to Eq. (7).

We next focus on $\|\tilde{r}_n\|_\infty$. Apply Hölder's inequality and triangle inequality:

$$\begin{aligned}\|\tilde{r}_n\|_\infty &\leq \sqrt{n}\|\hat{\Theta} - \Sigma_0^{-1}\|_{L_\infty}\left(\|\Delta_n\beta_0\|_\infty + \left\|\frac{1}{n}\tilde{X}^\top\varepsilon\right\|_\infty\right) \\ &\leq 2\sqrt{n}b_0b_1\tilde{\nu}_n\left(\|\Delta_n\beta_0\|_\infty + \left\|\frac{1}{n}\tilde{X}^\top\varepsilon\right\|_\infty\right).\end{aligned}$$

Here in the second line we invoke the conclusion in Lemma 3.1. It then suffices to upper bound $\|\Delta_n\beta_0\|_\infty$ and $\|\frac{1}{n}\tilde{X}^\top\varepsilon\|_\infty$. With Definition 6.1, it holds with probability $1 - o(1)$ that

$$\begin{aligned}\|\Delta_n\beta_0\|_\infty &\leq \left\|\left(\frac{1}{n}\tilde{X}^\top X - \Sigma_0\right)\beta_0\right\|_\infty + \left\|\left(\tilde{\Sigma} - \Sigma_0\right)\beta_0\right\|_\infty \\ &\leq \left[\varphi_{u,v}\left(\frac{1}{n}\tilde{X}^\top X, \Sigma_0; \log p\right) + \varphi_{u,v}\left(\tilde{\Sigma}, \Sigma_0; \log p\right)\right]\|\beta_0\|_2\end{aligned}$$

and

$$\left\|\frac{1}{n}\tilde{X}^\top\varepsilon\right\|_\infty \leq \sigma_\varepsilon\varphi_{\varepsilon,\infty}\left(\frac{1}{n}\tilde{X}\right).$$

By Lemmas 6.1 and 6.2, if $\frac{\log p}{\rho_*^4 n} \to 0$ then

$$\|\Delta_n\beta_0\|_\infty \leq O_{\mathbb{P}}\left\{\sigma_x^2\|\beta_0\|_2\sqrt{\frac{\log p}{\rho_*^2 n}}\right\} \quad \text{and} \quad \left\|\frac{1}{n}\tilde{X}^\top\varepsilon\right\|_\infty \leq O_{\mathbb{P}}\left\{\sigma_x\sigma_\varepsilon\sqrt{\frac{\log p}{\rho_* n}}\right\}.$$

E Tail inequalities

Lemma E.1 (Sub-Gaussian concentration inequality)**.** *Suppose X is a univariate sub-Gaussian random variable with parameter $\sigma > 0$; that is, $\mathbb{E}X = 0$ and $\mathbb{E}e^{tX} \leq e^{\sigma^2 t^2/2}$ for all $t \in \mathbb{R}$. Then*

$$\begin{aligned}\Pr\left[|X| \geq \epsilon\right] &\leq 2e^{-\frac{\epsilon^2}{2\sigma^2}}, \qquad \forall t > 0; \\ \mathbb{E}|X|^r &\leq r \cdot 2^{r/2} \cdot \sigma^r \cdot \Gamma\left(\frac{r}{2}\right), \qquad \forall r = 1, 2, \cdots\end{aligned}$$

Lemma E.2 (Sub-exponential concentration inequality)**.** *Suppose $X_1, \cdots, X_n$ are i.i.d. univariate sub-exponential random variables with parameter $\lambda > 0$; that is, $\mathbb{E}X_i = 0$ and $\mathbb{E}e^{tX_i} \leq e^{t^2\lambda^2/2}$ for all $|t| \leq 1/\lambda$. Then*

$$\Pr\left[\left|\frac{1}{n}\sum_{i=1}^n X_i\right| > \epsilon\right] \leq 2\exp\left\{-\frac{n}{2}\min\left(\frac{\epsilon^2}{\lambda^2}, \frac{\epsilon}{\lambda}\right)\right\}.$$

Lemma E.3 (Hoeffding inequality)**.** *Suppose $X_1, \cdots, X_n$ are independent univariate random variables with $X_i \in [a_i, b_i]$ almost surely. Then for all $t > 0$, we have that*

$$\Pr\left[\left|\frac{1}{n}\sum_{i=1}^n X_i - \mathbb{E}X_i\right| > t\right] \leq 2\exp\left\{-\frac{2n^2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right\}.$$

Lemma E.4 (Bernstein inequality, support condition)**.** *Suppose $X_1, \cdots, X_n$ are independent random variables with zero mean and finite variance. If $|X_i| \leq M < \infty$ almost surely for all $i = 1, \cdots, n$, then*

$$\Pr\left[\left|\frac{1}{n}\sum_{i=1}^{n} X_i\right| > t\right] \leq 2\exp\left\{-\frac{\frac{1}{2}n^2t^2}{\sum_{i=1}^{n}\mathbb{E}X_i^2 + \frac{1}{3}Mnt}\right\}, \qquad \forall t > 0.$$

Lemma E.5 (Bernstein inequality, moment condition)**.** *Suppose $X_1, \cdots, X_n$ are independent random variables with zero mean and $\mathbb{E}|X_i|^2 \leq \sigma^2 < \infty$. Assume in addition that there exists some positive number $L > 0$ such that*

$$\mathbb{E}|X_i|^k \leq \frac{1}{2}\sigma^2 L^{k-2}k!, \qquad \forall k > 1.$$

Then we have that

$$\Pr\left[\left|\frac{1}{n}\sum_{i=1}^{n} X_i\right| > t\right] \leq 2\exp\left\{-\frac{nt^2}{2(\sigma^2 + Lt)}\right\}, \qquad \forall t > 0.$$

Lemma E.6 (Hsu et al. (2012))**.** *Suppose $X = (X_1, \cdots, X_p)$ is a p-dimensional zero-mean sub-Gaussian random vector; that is, there exists $\sigma > 0$ such that*

$$\mathbb{E}\exp\left\{\alpha^\top X\right\} \leq \exp\left\{\|\alpha\|_2^2\sigma^2/2\right\}, \qquad \forall \alpha \in \mathbb{R}^p.$$

Let A be a $p \times p$ positive semi-definite matrix. Then for all $t > 0$,

$$\Pr\left[X^\top AX > \sigma^2\left(\mathrm{tr}(A) + 2\sqrt{\mathrm{tr}(A^2)t} + 2\|A\|_{\mathrm{op}}t\right)\right] \leq e^{-t}.$$

References

Cai, T., Liu, W., & Luo, X. (2011). A constrained L1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association, 106*(494), 594–607.

Hsu, D., Kakade, S. M., & Zhang, T. (2012). A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability, 17*(52), 1–6.

Loh, P.-L., & Wainwright, M. (2012). High-dimensional regression with noisy and missing data: provable guarantees with nonconvexity. *The Annals of Statistics, 40*(3), 1637–1664.

van de Geer, S. (2010). *Empirical Processes in M-Estimation*. Cambridge University Press.